wax
13 FEB 1975

PRE-PUBLICATION DRAFT

A SIMPLE MIXED-TRAFFIC TECHNIQUE FOR GROUND-BASED ALOHA CHANNELS

by

Richard Binder
University of Hawaii
Honolulu, Hawaii

January 1975

Sponsored by
Advanced Research Projects Agency
ARPA Order No. 1956

*****

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency of the United States Government.

A SIMPLE MIXED-TRAFFIC TECHNIQUE FOR GROUND-BASED ALOHA CHANNELS

by

Richard Binder
University of Hawaii

Abstract

This report concerns the problem of how to regulate file traffic in a slotted or unslotted ALOHA channel when interactive user traffic is also present. Some simulation results are given for a simple technique in which a constant interval of time is inserted by all file nodes between successive new packet transmissions. The interactive traffic nodes operate in a normal (uncontrolled) manner; the constant time interval used by file nodes is adjusted to achieve the desired delay characteristics for the interactive users. The technique is intended for a ground-based environment in which interactive user traffic predominates, with only occasional requirements for file transfers. It does not appear efficient for satellite channels due to their large propagation delays.

This report was supported by THE ALOHA SYSTEM, a research project at the University of Hawaii, which is supported by the Advanced Research Projects Agency of the Department of Defense and monitored by NASA Ames Research Center under Contract No. NAS2-8590.

ACKNOWLEDGEMENT

The author wishes to thank Rodney Okano, Neil Iwamoto, Lorraine Sakaguchi, and Peter Law for their simulator programming, and Michael Ferguson for several helpful discussions during the development of the results.

Digitized by the Internet Archive
in 2025 with funding from
Amateur Radio Digital Communications, Grant 151

https://archive.org/details/aloha-05

A SIMPLE MIXED-TRAFFIC TECHNIQUE FOR GROUND-BASED ALOHA CHANNELS

Table of Contents

		Page
	ABSTRACT	i
	ACKNOWLEDGEMENT	ii
	LIST OF FIGURES AND TABLES	iv
1.	INTRODUCTION	1
2.	THE TECHNIQUE	2
	File Users Only	2
	The General Case	3
	Implementation	4
	Discussion	4
3.	RESULTS WITH FILE USERS ONLY	6
4.	MIXED TRAFFIC RESULTS	13
5.	CONCLUSIONS	23
	REFERENCES	24
	APPENDIX I: SIMULATION DETAILS	25

List of Figures and Tables

Figure		Page
1	OPTIMIZED SYNC TIME VS FRAME RATIO	7
2	THRUPUT PRIOR TO SYNC	8
3	TOTAL FILE THRUPUT VS NUMBER OF FILE USERS	9
4	FILE DELAY VS FRAME RATIO	12
5	MIXED TRAFFIC PERFORMANCE VS FRAME RATIO (F=10,I=10)	14
6	I-USER DELAY VS FILE THRUPUT, 3 FILE USERS	15
7	I-USER DELAY VS FILE THRUPUT, 10 FILE USERS	16
8	FILE THRUPUT WITH CONSTANT I-DELAY, 3 FILE USERS	17
9	FILE THRUPUT WITH CONSTANT I-DELAY, 10 FILE USERS	18
10	FILE THRUPUT VS FRAME RATIO, 3 FILE USERS	20
11	FILE THRUPUT VS FRAME RATIO, 10 FILE USERS	21

Table		
1	OPTIMUM R_e VALUES FOR FIGURE 1	10
2	OPTIMUM R_e VALUES FOR MIXED TRAFFIC RESULTS	22

1. INTRODUCTION

The random access (ALOHA channel) multiplexing technique is based on a traffic environment characterized by a relatively large number of users generating high peak-to-average, or bursty, traffic. In a practical implementation such as the ALOHANET, however, a need occasionally arises to transfer a file from a user node (perhaps stored on paper tape, cassette, or in a minicomputer) to another node over the random access channel. If the file is transmitted without constraints, serious degradation of service can result to the interactive users in the form of large increases in packet delay. Also, if two or more nodes are sending files at the same time, the result may be a relatively small thruput for all users.

A previous analysis [1] has investigated the effects of allowing one or more 'large' users to send randomly-generated packets with a large mean rate compared to the population of interactive user traffic. While relatively high thruput values can be achieved ("excess capacity"), the interactive user traffic suffers significantly larger delays than when the large users are not present [2]. Reservation techniques [3,4,5] have been designed to allow more efficient channel use with respect to both thruput and delay in the context of a satellite channel and relatively concentrated node traffic. While these techniques appear to be also applicable to ground-based contexts such as the ALOHANET, their implementation costs appear high relative to the gains expected when interactive traffic predominates.

This report describes a relatively simple technique in which a single parameter is used to control all nodes sending file traffic, with the parameter dynamically adjusted as both the number of file users and interactive users change. The method allows any desired delay to be achieved for the interactive users (relative to a channel with no file traffic), and also reduces or eliminates interferences between the file users.

The technique is defined in detail in the next section for an environment consisting of either a slotted or unslotted channel without carrier-sense. Simulation results are then given which characterize the file traffic performance when no interactive user traffic is present. Following this, results are presented for some mixtures of file and interactive user traffic.

2. THE TECHNIQUE

We first discuss the case when only file users are present. Interactive users are then introduced and the general control method defined.

File Users Only

Assume for the moment a slotted channel with exactly two users. Suppose a random retransmission discipline is used by both users with the same mean time R_e (expressed in multiples of a packet transmission time, or slot). Suppose further that both users wish to send a long file, and that each can send the file continuously at the channel data rate if allowed (i.e., in a stream of packets with no separation in time). If they both begin sending at the same time, they will conflict and randomly retransmit. As soon as one is successful, that user immediately sends another new packet. The result is that they continue to conflict with each other, with a thruput dependent on the value of R_e used.

Suppose now that, instead of sending a new packet immediately after a previous success, each user waits a constant interval after the success before sending the next new packet. Let this interval be exactly one slot for each user. Now the following will happen: after some finite time following their first transmissions, the retransmissions of each user will 'land' in consecutive slots. Following these successes each user will allow one slot to go by before sending a new packet, resulting in an alternation of new packet transmissions by each user without further conflicts.

Let us now generalize the preceding example by defining the constant interval in terms of a frame ratio FR, such that if N users each allow S slots to go by following a success before sending their next new packet,

$$FR = \frac{N}{1+S} \qquad (1)$$

where the term 1+S includes the successful packet transmission. Let us also extend the notion to an unslotted channel by expressing S as a number (not necessarily an integer) of packet transmission times, where the latter is

equal to a slot time for the slotted case. We shall restrict FR to be between zero and one.

The previous result for two users follows for any number of users N: once they are all successful, there will be no further conflicts (since all use the same constant interval S). Let us refer to this conflict-free condition as "sync". The time it takes for this to happen will depend on the value of FR for a given N -- the smaller FR is made (by making S larger), the higher the probability of sync occurring within a given time. We will call this interval the sync time T_S, defined as the mean time elapsing between the first (simultaneous) transmissions by all users and the last retransmission by the users.

Let the total channel thruput (or just "thruput") be the fraction of time the channel is occupied by successfully sent data plus overhead bits. Following the occurrence of sync, the thruput is equal to the value of FR (and equals one for the example with two users). There will also be a non-zero thruput value averaged over the initial interval T_S; let us call this value S_T. Thus for a given file length (assumed the same for all users), the total average thruput will be determined by the values of T_S, S_T, and FR (call this total the "file thruput"). In general, there will be an optimal value of FR for a given file length and number of users which maximizes the file thruput. Note that for a large enough N or small enough file length, this optimum may not result in sync prior to the end of the file transmissions.

The General Case

Assume now that interactive users (I-users) are also present, and that each generates a new packet following a previous success after a random think-time interval ("IPT"). Let these I-users use the same retransmission discipline (with the same mean R_e) as the file users.

To achieve the desired average I-user packet delay, we simply decrease the value of FR used by all file users, trading file thruput for improved I-user delay. In the limit as FR is made zero, I-user delay is equal to the value achievable when no file traffic is present and file thruput is zero. Note that the file users may frequently be removed from a sync condition following an I-user transmission, depending on the FR value being used. This will also

occur in the more realistic case where file users come and go, rather than all starting at the same time as defined above.

We may also remove the constraint originally assumed which required all file users to be able to send their file at the channel data rate: it is only necessary that they send a new packet at integer multiples of 1+S. In fact, for channels with reasonable numbers of active I-users and/or file users S will be large, allowing only a fraction of the channel bandwidth to be used by each file user (reducing the need to wait longer than the first S).

Implementation

The implementation consists simply of providing the currently desired value of S in (1) to all file user nodes, which adjust their new packet transmissions accordingly. The determination of S is a function of the active number of I-users and file users, and the desired I-user delay. It could easily be generated by a centralized node such as the ALOHANET's MENEHUNE and broadcast whenever necessary to all user nodes. The MENEHUNE already maintains table information on all active users; those wishing to send a file could first transmit a special information packet to the MENEHUNE, providing the latter with all necessary information for determining S. A distributed implementation method is of course also possible in which each file node identically determines the S value -- this may, however, be much costlier to achieve.

The next sections will present simulation data which could be used in table form to determine S (with the simulation data expanded to cover the desired range of operation). Alternatively, an analytical model might be developed which will allow efficient calculation of S.

Discussion

Readers familiar with the satellite reservation schemes will have noted some similarities between the frame ratio technique just presented and the frame concept of Reservation ALOHA [3]. The latter also utilizes a frame interval which elapses between successive uses of a particular slot by each

user. Beyond this, however, there are some significant differences. In Reservation ALOHA (R-ALOHA) each user can use many slots of a frame, rather than just one. Furthermore, once successful a user is guaranteed use of a slot without further interference (barring local noise problems) -- this is not true in the FR scheme, in which a user may be subsequently conflicted with by retransmissions from other file users or by new packets from both I-users and new file users. Both factors will produce significantly different performance results, depending on the assumed traffic characteristics.*

Moreover, the two schemes have been developed for quite different contexts. R-ALOHA (and other satellite reservation schemes) attempt to dynamically allocate a channel to nodes with concentrated traffic while minimizing the delays introduced by the long satellite propagation time. The FR scheme attempts to regulate a small number of file users so as not to degrade service to a larger population of high peak-to-average ratio I-users, while also providing good thruput characteristics for the file users when the channel is lightly loaded by I-users.

This is not to say that one or more of the satellite reservation schemes might not provide more cost-effective performance than the FR scheme in the latter's context. The effort reported in this paper has been simply to investigate the properties of a simple scheme which does not require control of the I-users, and which is limited to a ground-based channel.

---

* N.T. Gaarder has pointed out that a generalized model of R-ALOHA developed by E. Kanehira [6] also models the FR scheme as a special case, if no I-users are allowed and multiple transmissions by each user within a frame are suppressed in the model.

3. RESULTS WITH FILE USERS ONLY

Simulations were performed with the simulator described in [7]; the relevant simulation details are given in Appendix I. The first step was to quantify the sync times T_S for different values of N and FR with no I-users present. The thruput prior to sync, S_T, was also measured for each T_S. Given this data, the resulting optimal file thruput was then calculated for different values of file lengths and numbers of file users.

Figure 1 shows the resulting T_S values as a function of FR for constant numbers of file users, for both slotted and unslotted channels. The units of T_S are given in 720-bit packet transmission times, equal to a slot time. Each data point represents an average of approximately eight independent simulation runs (each with a different starting seed). Each of these averaged points was optimized with respect to the retransmission mean R_e -- that is, a separate set of runs was made for each of several values of R_e, with the set selected which had the smallest average value of T_S. The corresponding optimal values of R_e are given in Table 1. (Note that R_e was limited to integer values by the simulator.)

Figure 2 shows the values of S_T corresponding to the sync times of Figure 1. Note that fairly high values are obtained during this period relative to the maximum 1/e or 1/2e expected in an ALOHA channel with Poisson traffic. The thruput values following the occurrence of sync are given directly by the corresponding FR values.

Figure 3 shows the resulting file thruput as a function of the number of file users, for constant file lengths. Each point in Figure 3 represents the maximum file thruput achievable from the set of all FR values tried for a given number of users. If sync occurred prior to the end of the file transmissions for a value of FR, the following equation was used (where S_F represents total file thruput):

$$S_F = \frac{N \cdot L}{T_S + (N \cdot L - G)/FR} \qquad (2)$$

where N is the number of file users,
L is the file length in packets sent by each user,
G is the number of packets sent successfully by all users prior to sync, and is given by

Figure 1. OPTIMIZED SYNC TIME VS FRAME RATIO

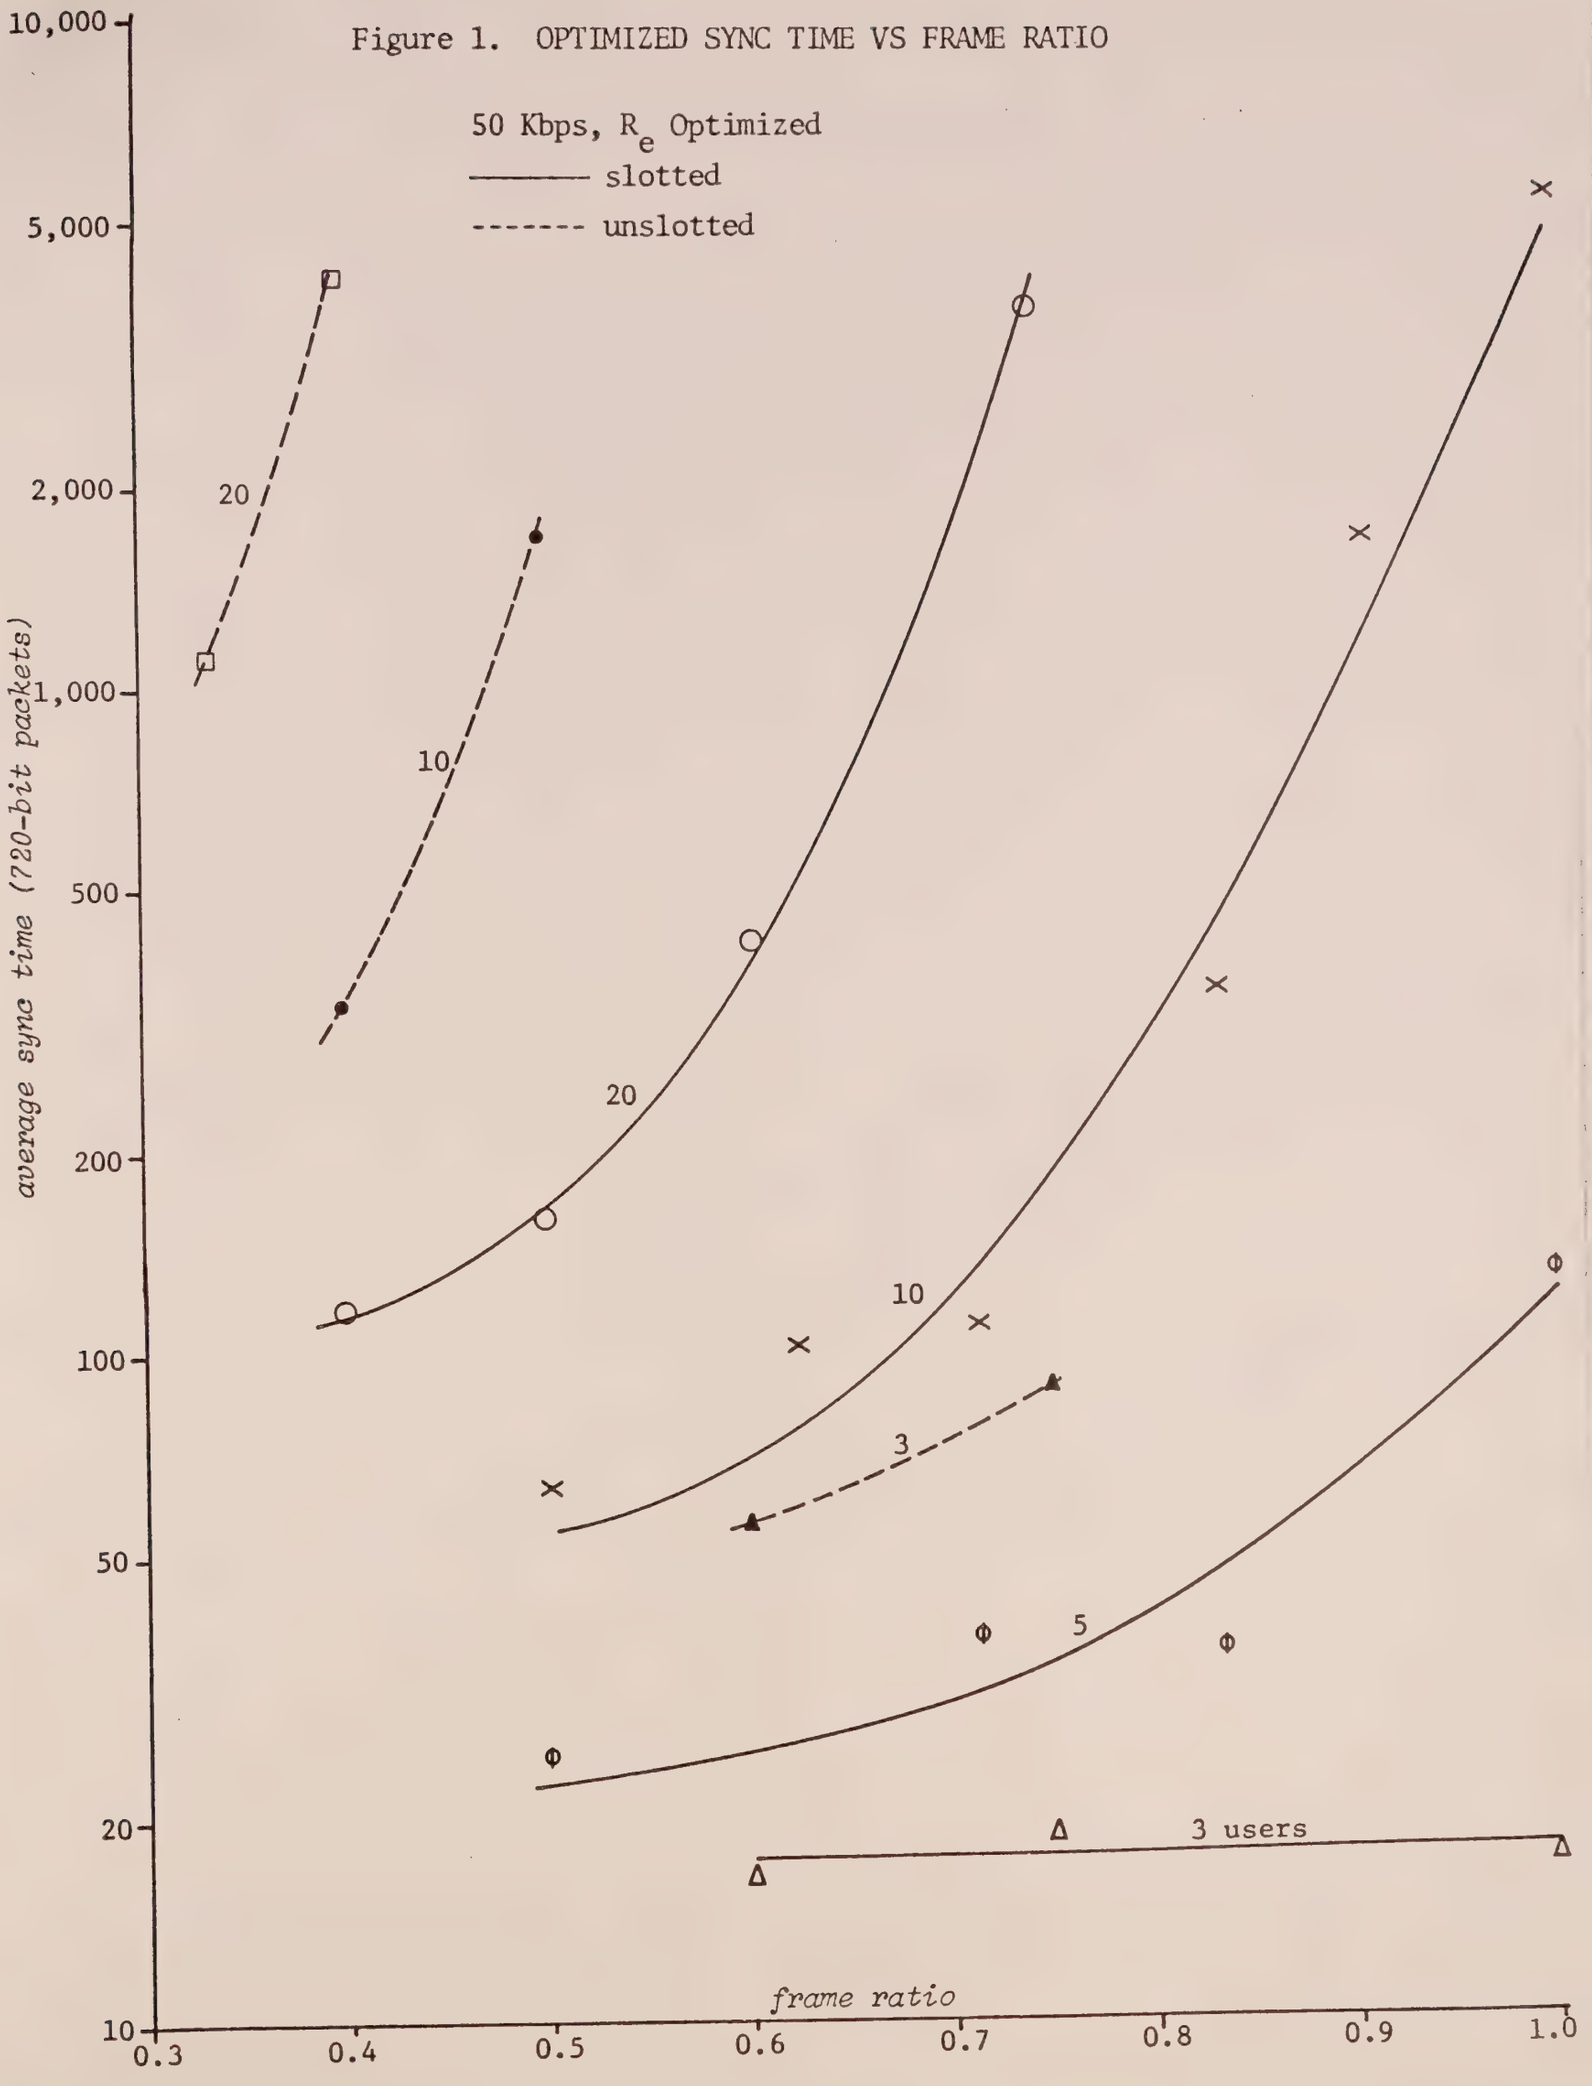

Figure 2. THRUPUT PRIOR TO SYNC

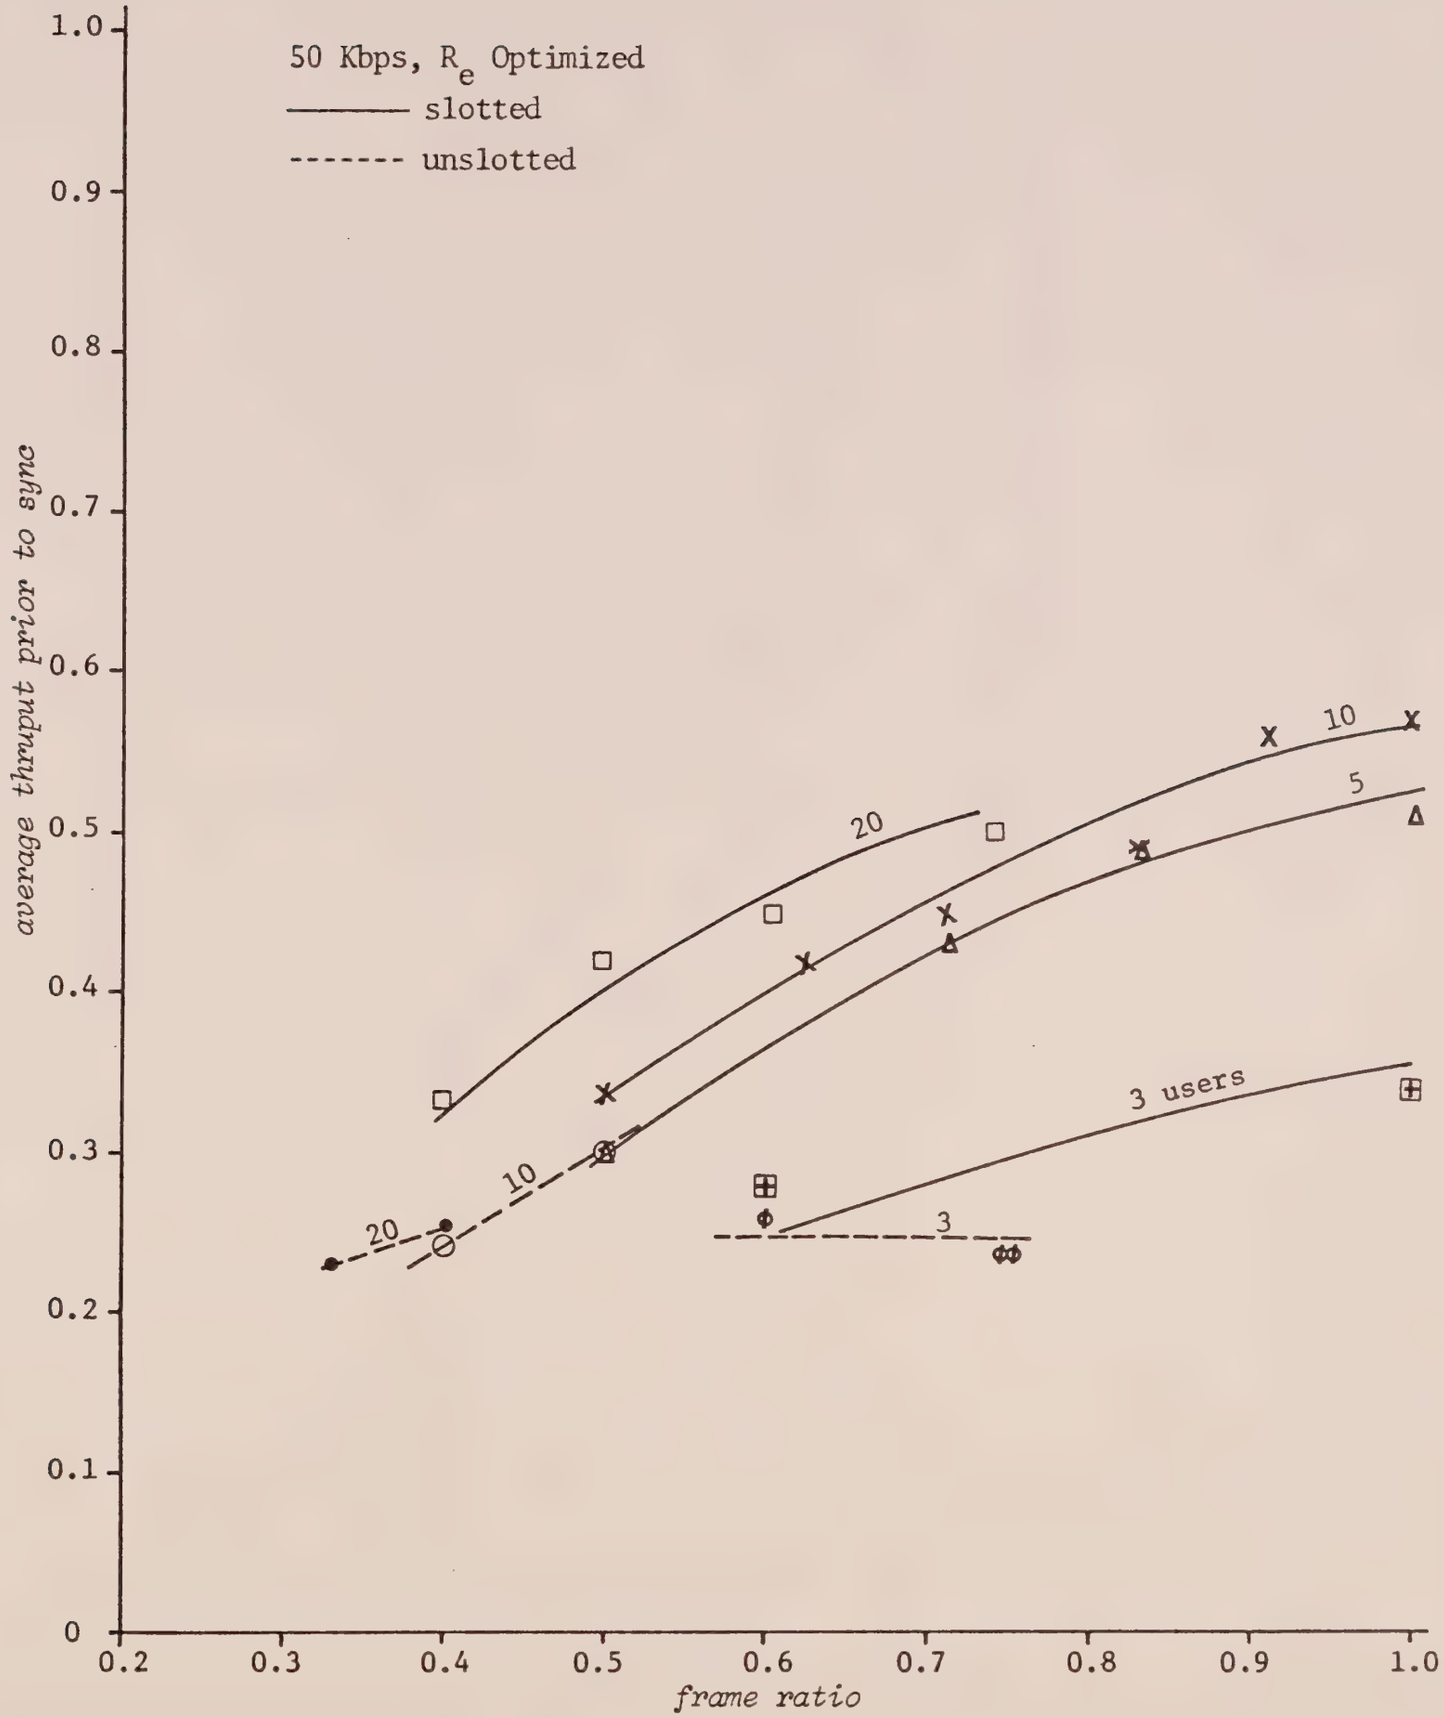

Figure 3. TOTAL FILE THRUPUT VS NUMBER OF FILE USERS

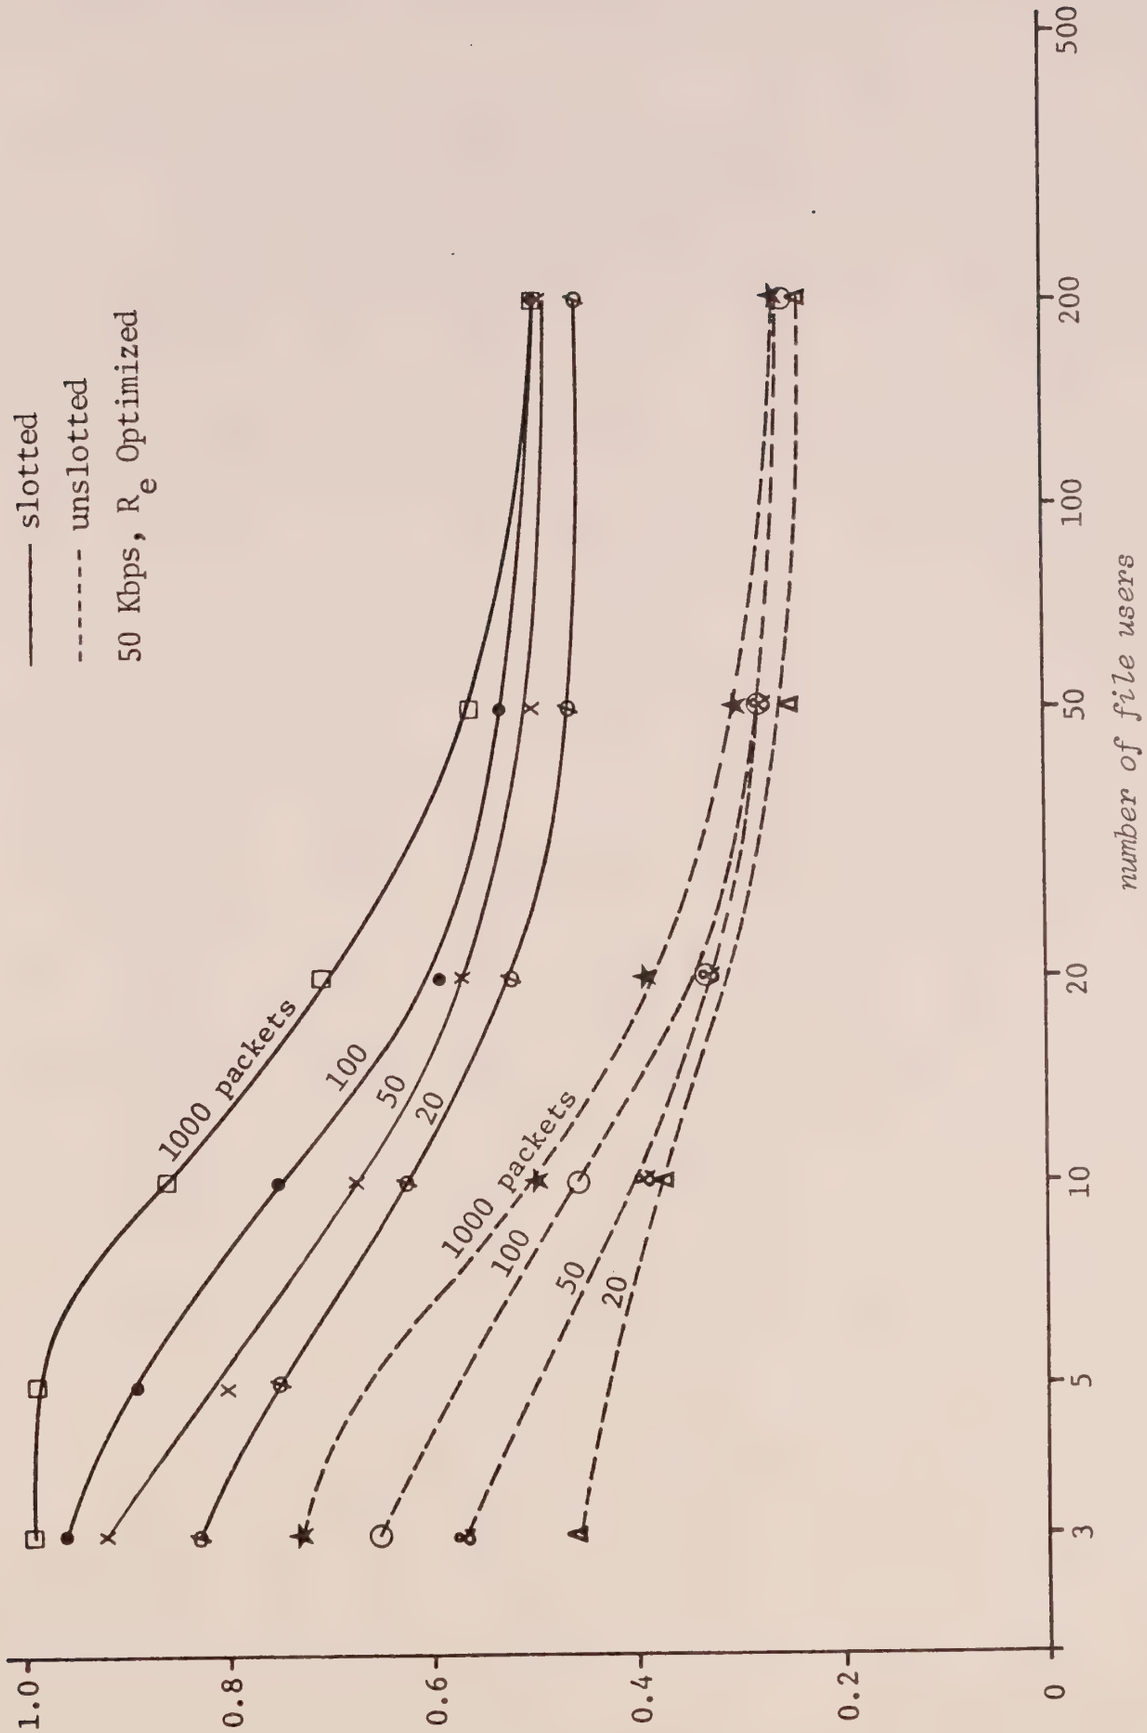

Table 1. OPTIMUM R_e VALUES FOR FIGURE 1

N	FR	Optimum R_e (720-bit Packets)	
		slotted	*unslotted*
3	3/3	1	20
	3/4	2	2
	3/5	2	2
5	5/5	2	
	5/6	2	(no runs
	5/7	2	made)
	5/10	2	
10	10/10	20	-
	10/11	15	-
	10/12	5	-
	10/14	10	-
	10/16	5	-
	10/20	4	15
	10/25	-	10
20	20/27	20	-
	20/33	10	-
	20/40	10	-
	20/50	5	30
	20/60	-	20

$$G = S_T \cdot T_S \tag{3}$$

In words, S_F is the ratio of total file packets sent to the total number of slots (or packet times) required to send them. If sync did not occur by the end of the files, the average thruput achieved at that time was used for the corresponding FR.

To give an indication of the degradation in file performance if non-optimal values of FR are used, the average file delay is shown in Figure 4 for 10 users as a function of FR. The delay is defined as the total elapsed time required to send the file by each user, and is obtained from the file thruput values by Equation (4):

$$D_F = \frac{N \cdot L \cdot 720}{R \cdot S_F} \tag{4}$$

where R is the channel data rate in bits per second and the factor of 720 converts the file length L to bits.

Note in Figure 4 that a relatively small change in elapsed time occurs over the range of FR values. In particular, the curve for an L of 100 ranges from 19 seconds to 23 seconds, which would probably not be considered significant by most users. The range for 1000 packets is from a little under 3 minutes to almost 4 minutes, a time difference which could generate impatience. Overall, however, the results suggest that a single FR value (for the given number of users) could be used for all the file lengths, eliminating a need for knowledge of the latter in a practical implementation.

Figure 4. FILE DELAY VS FRAME RATIO

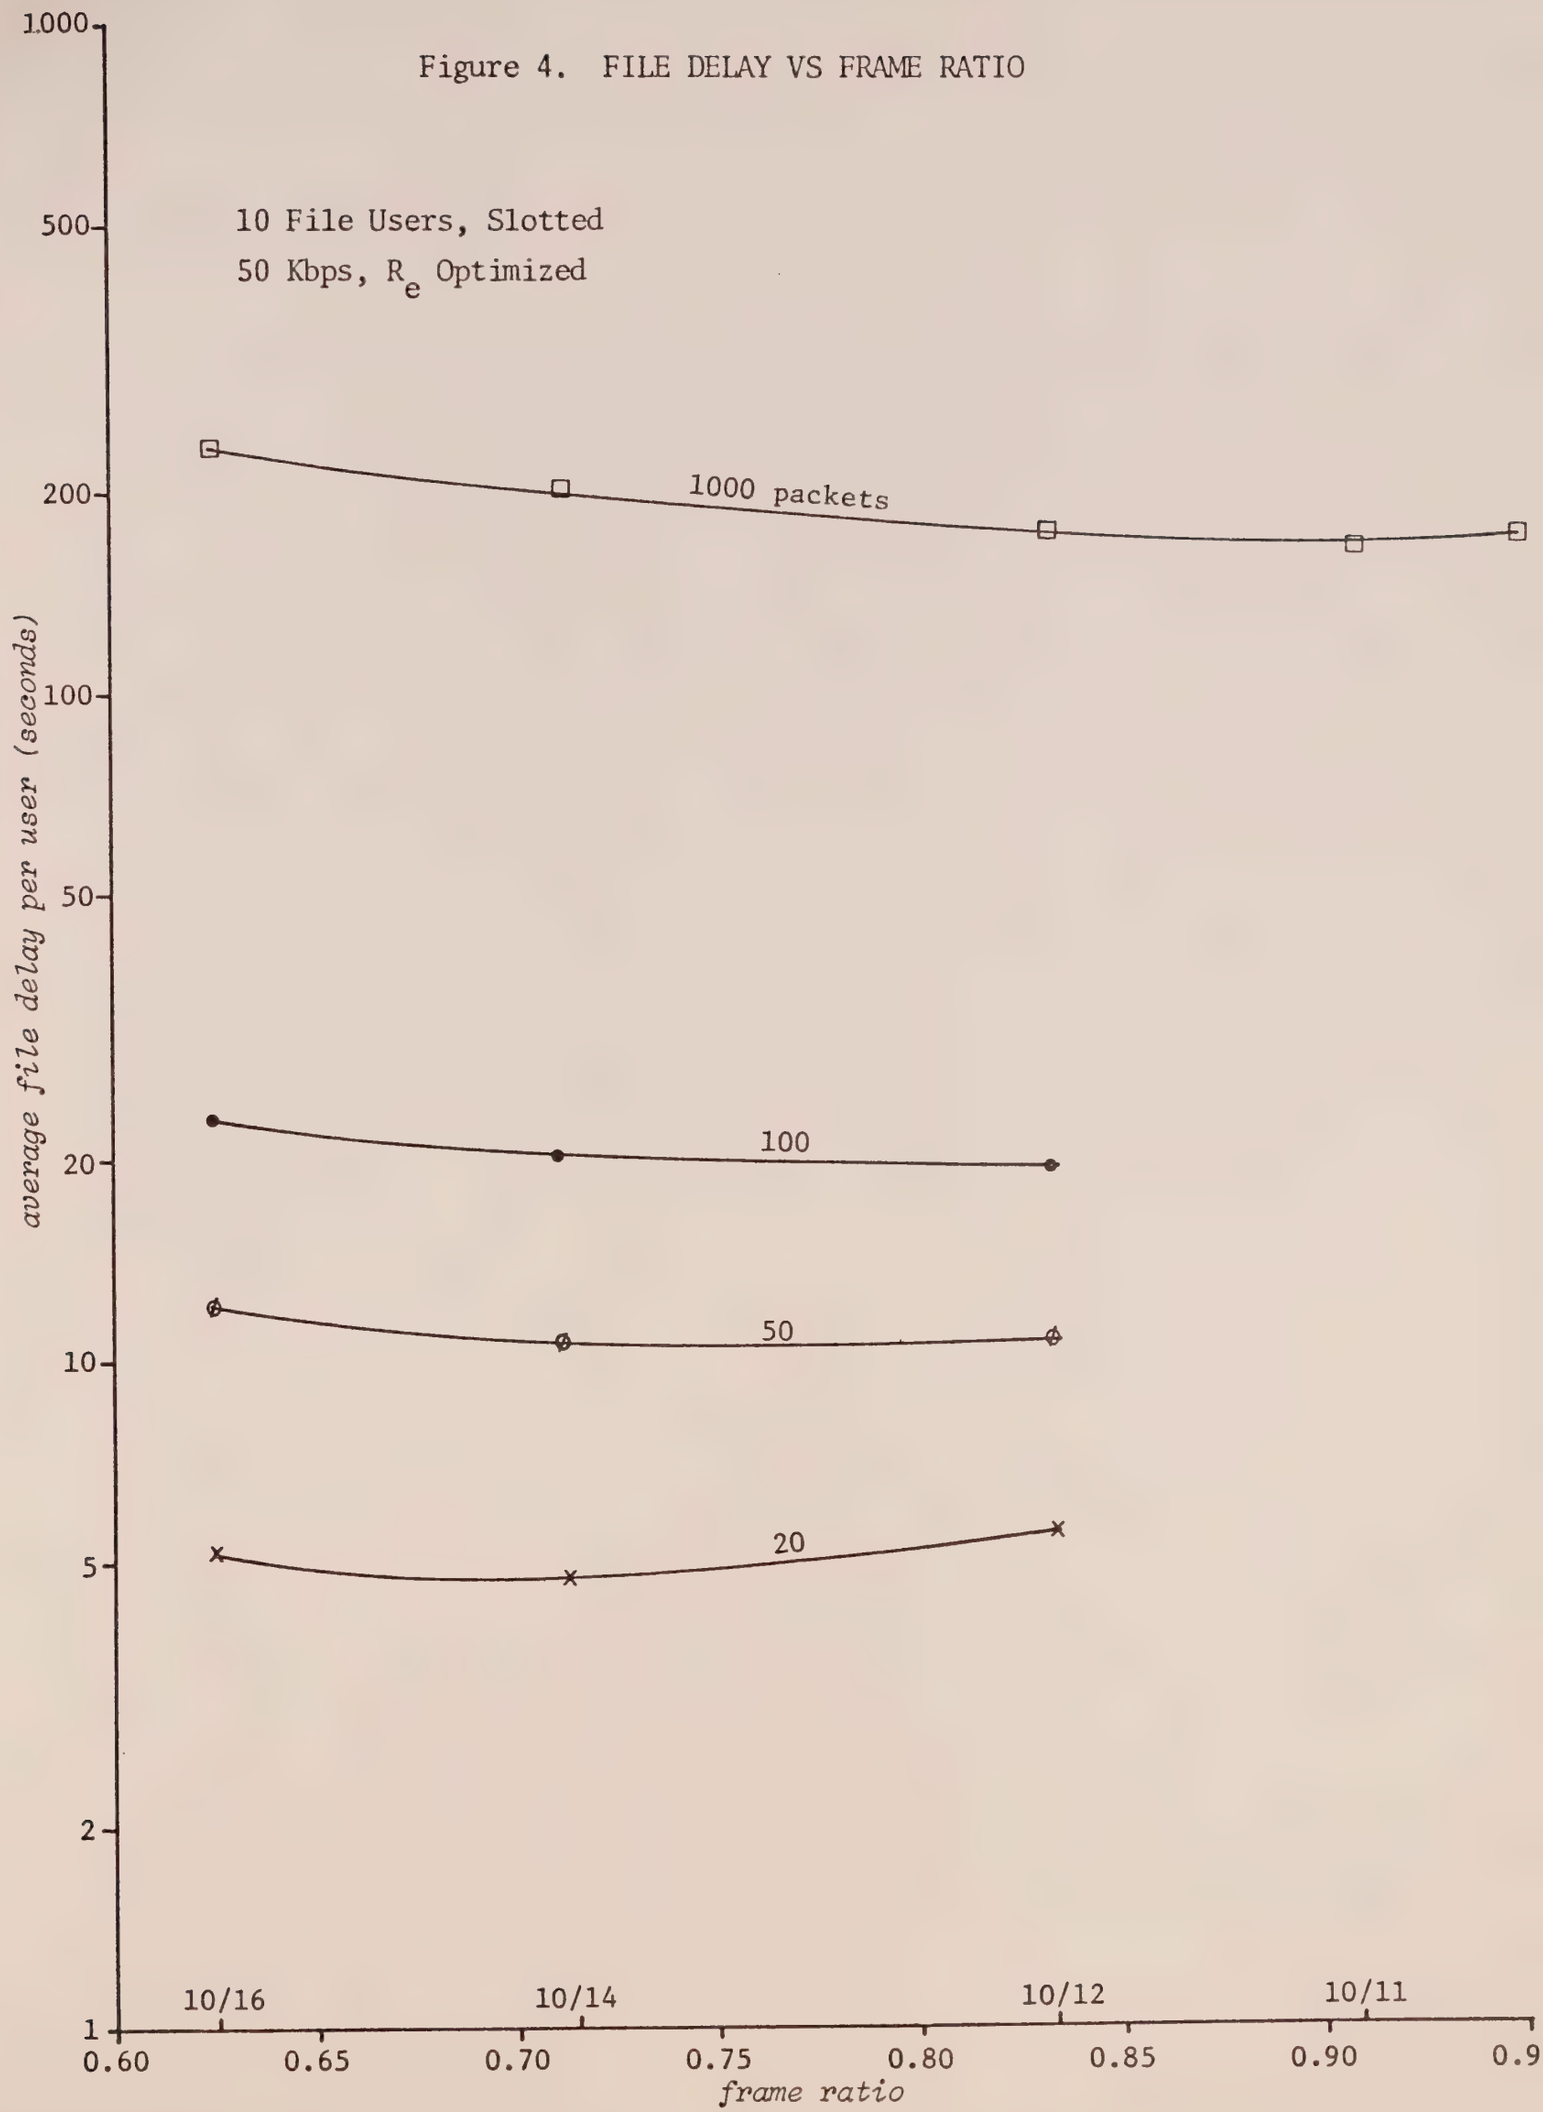

4. MIXED TRAFFIC RESULTS

In this section we present some simulation results for the case in which a given number of I-users are present along with the file users. The measurement approach is somewhat different, since now file users are forced to repeatedly re-sync as I-user transmissions disturb their use of the channel. Thus the average file-user thruput was measured for each simulation run, the latter consisting of a total of 10,000 packets (both new and retransmissions) generated by all users. This file thruput can be thought of as resulting from a population of file users continuously entering and leaving the system, such that their total number is constant. No specific file length is assumed. The average I-user packet delay was also obtained from each run. A mean think-time of 30 seconds was used for each I-user, with all users generating constant-length 720-bit packets. All results are for a 50 Kbps data rate.

For a given number of file users, I-users and a frame ratio, runs were made with different values of R_e to find the optimal R_e (defined as that which maximized file-user thruput). The resulting file-user thruput and I-user delay for optimal R_e are shown in Figure 5 as a function of FR for 10 file users and 10 I-users in a slotted channel. As can be seen, increasing FR causes a corresponding increase in I-user delay; above a certain value of FR, file thruput begins to decrease as well.

The results typified by Figure 5 are shown more generally in Figures 6 and 7, where we now plot average I-user delay versus file thruput. The independent variable being manipulated is still the frame ratio FR, where increasing values of FR correspond to moving upwards along each curve.

Figure 6 shows the results for 3 file users and either 10 or 50 I-users, for both a slotted and unslotted channel. The total I-user thruput S_I is also shown for each curve. Note that for the unslotted cases, as FR is increased I-user delay increases and file thruput decreases, corresponding to the familiar channel overload condition of a homogeneous population ALOHA channel. Figure 7 shows similar results for 10 file users and 10, 50 and 100 I-users.

Figures 8 and 9 show the resulting file thruput as a function of the number of I-users for constant I-user delay, for 3 and 10 file users respectively. Note the large values of file thruput achieved for 3 file users when 50 I-users

Figure 5. MIXED TRAFFIC PERFORMANCE VS FRAME RATIO (F=10, I=10)

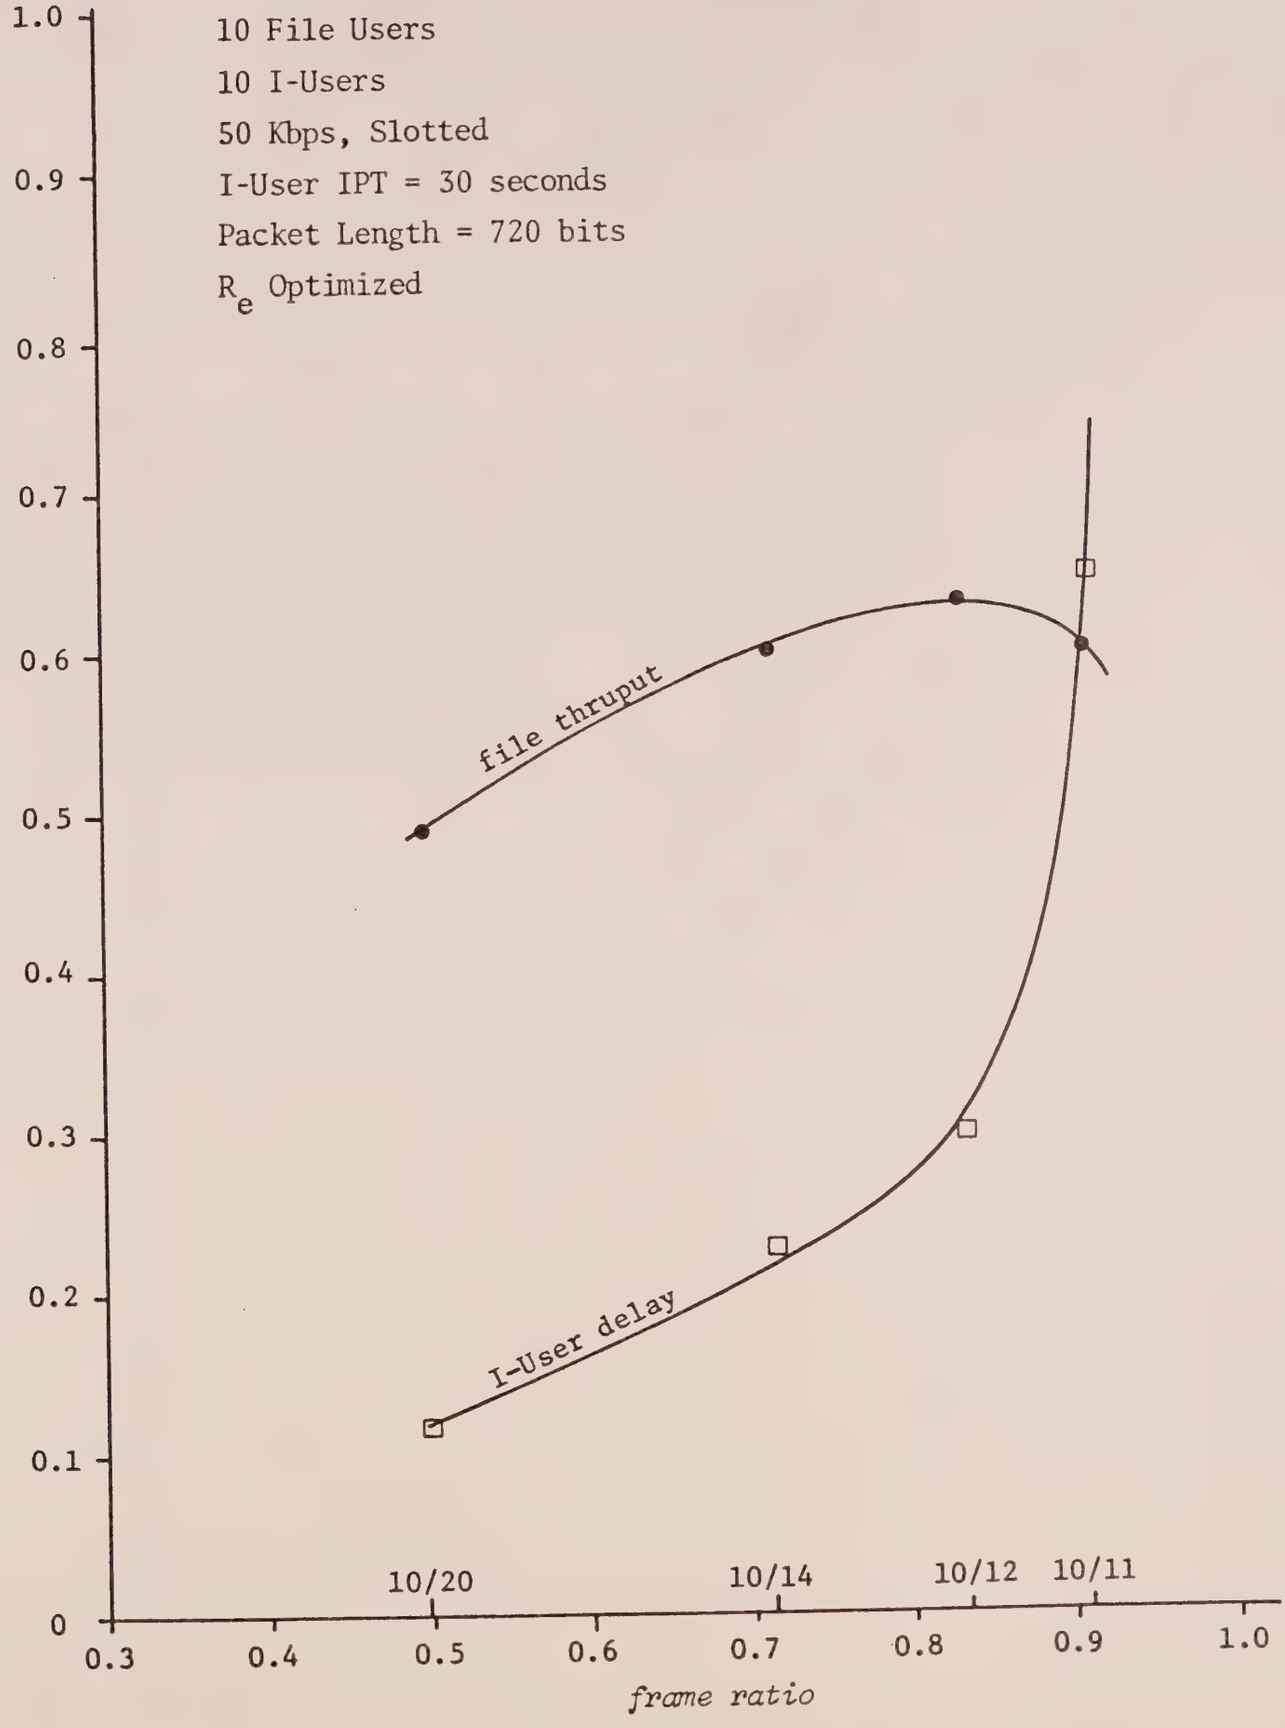

Figure 6. I-USER DELAY VS FILE THRUPUT, 3 FILE USERS

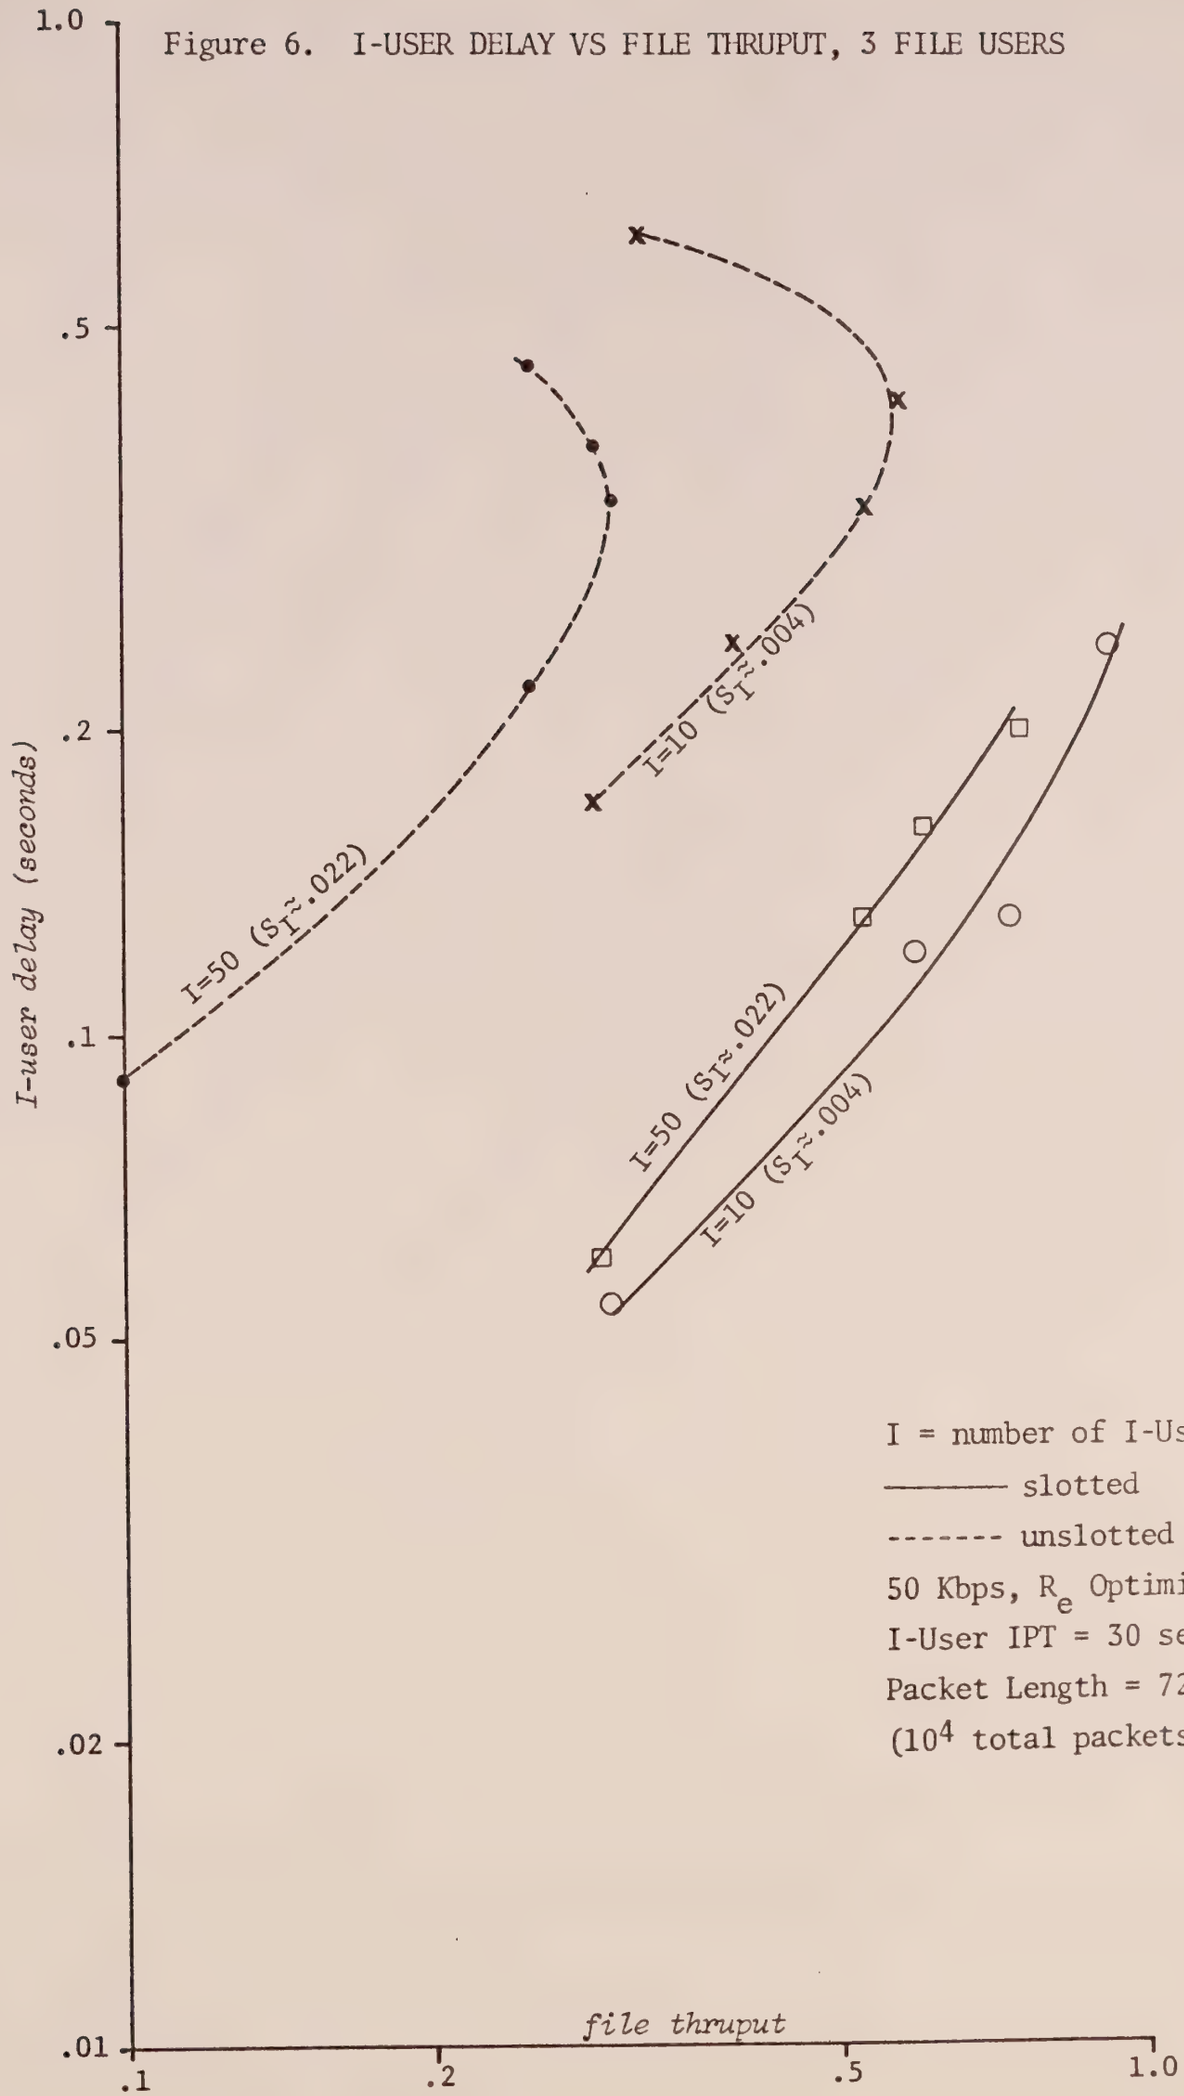

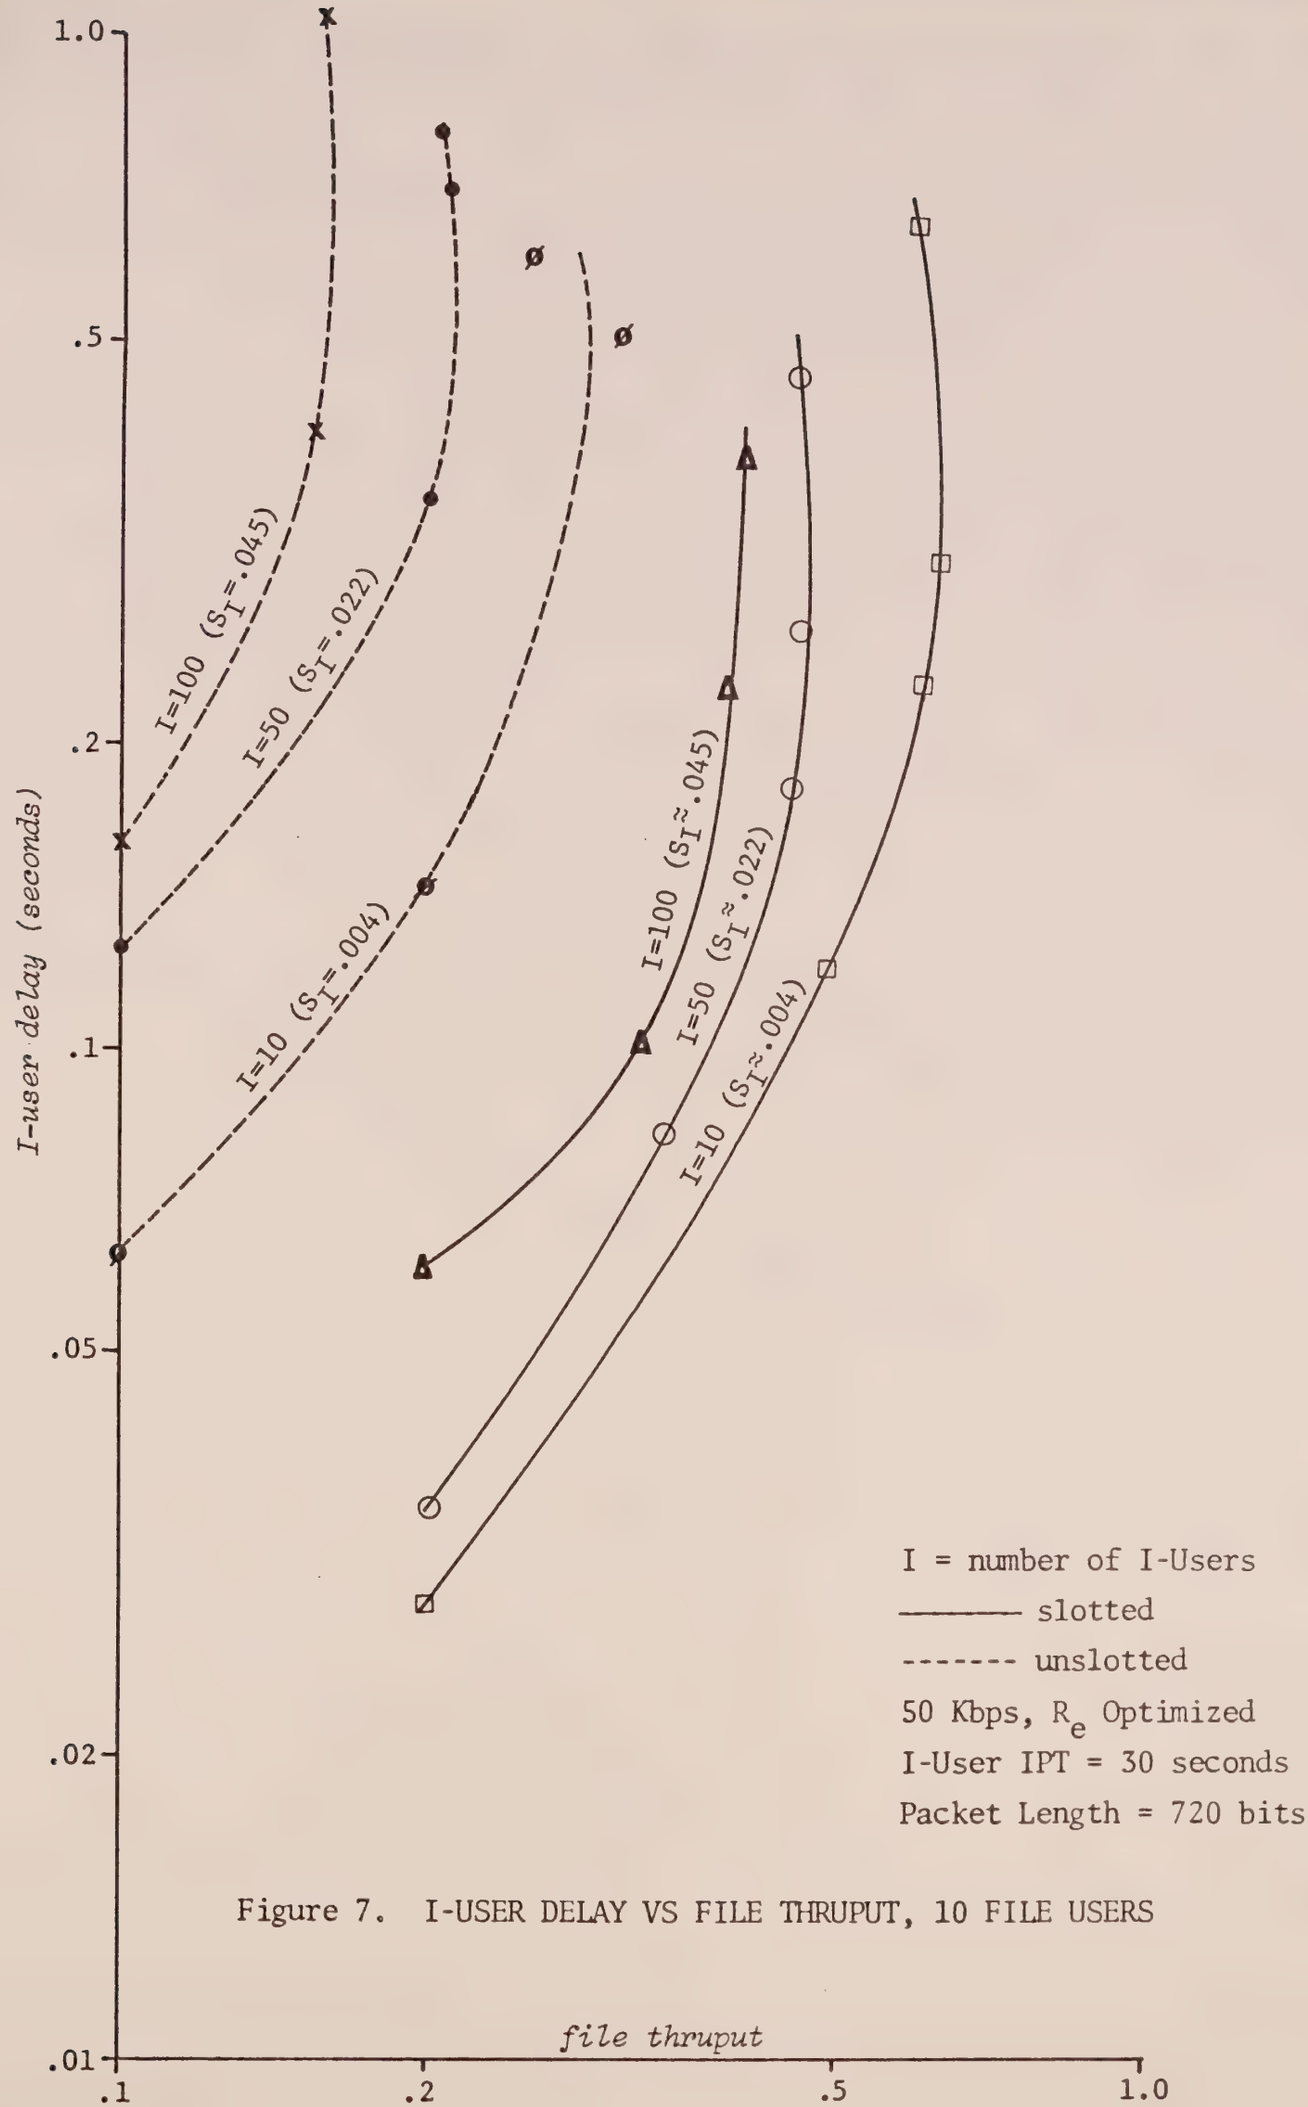

Figure 7. I-USER DELAY VS FILE THRUPUT, 10 FILE USERS

Figure 8. FILE THRUPUT WITH CONSTANT I-DELAY, 3 FILE USERS

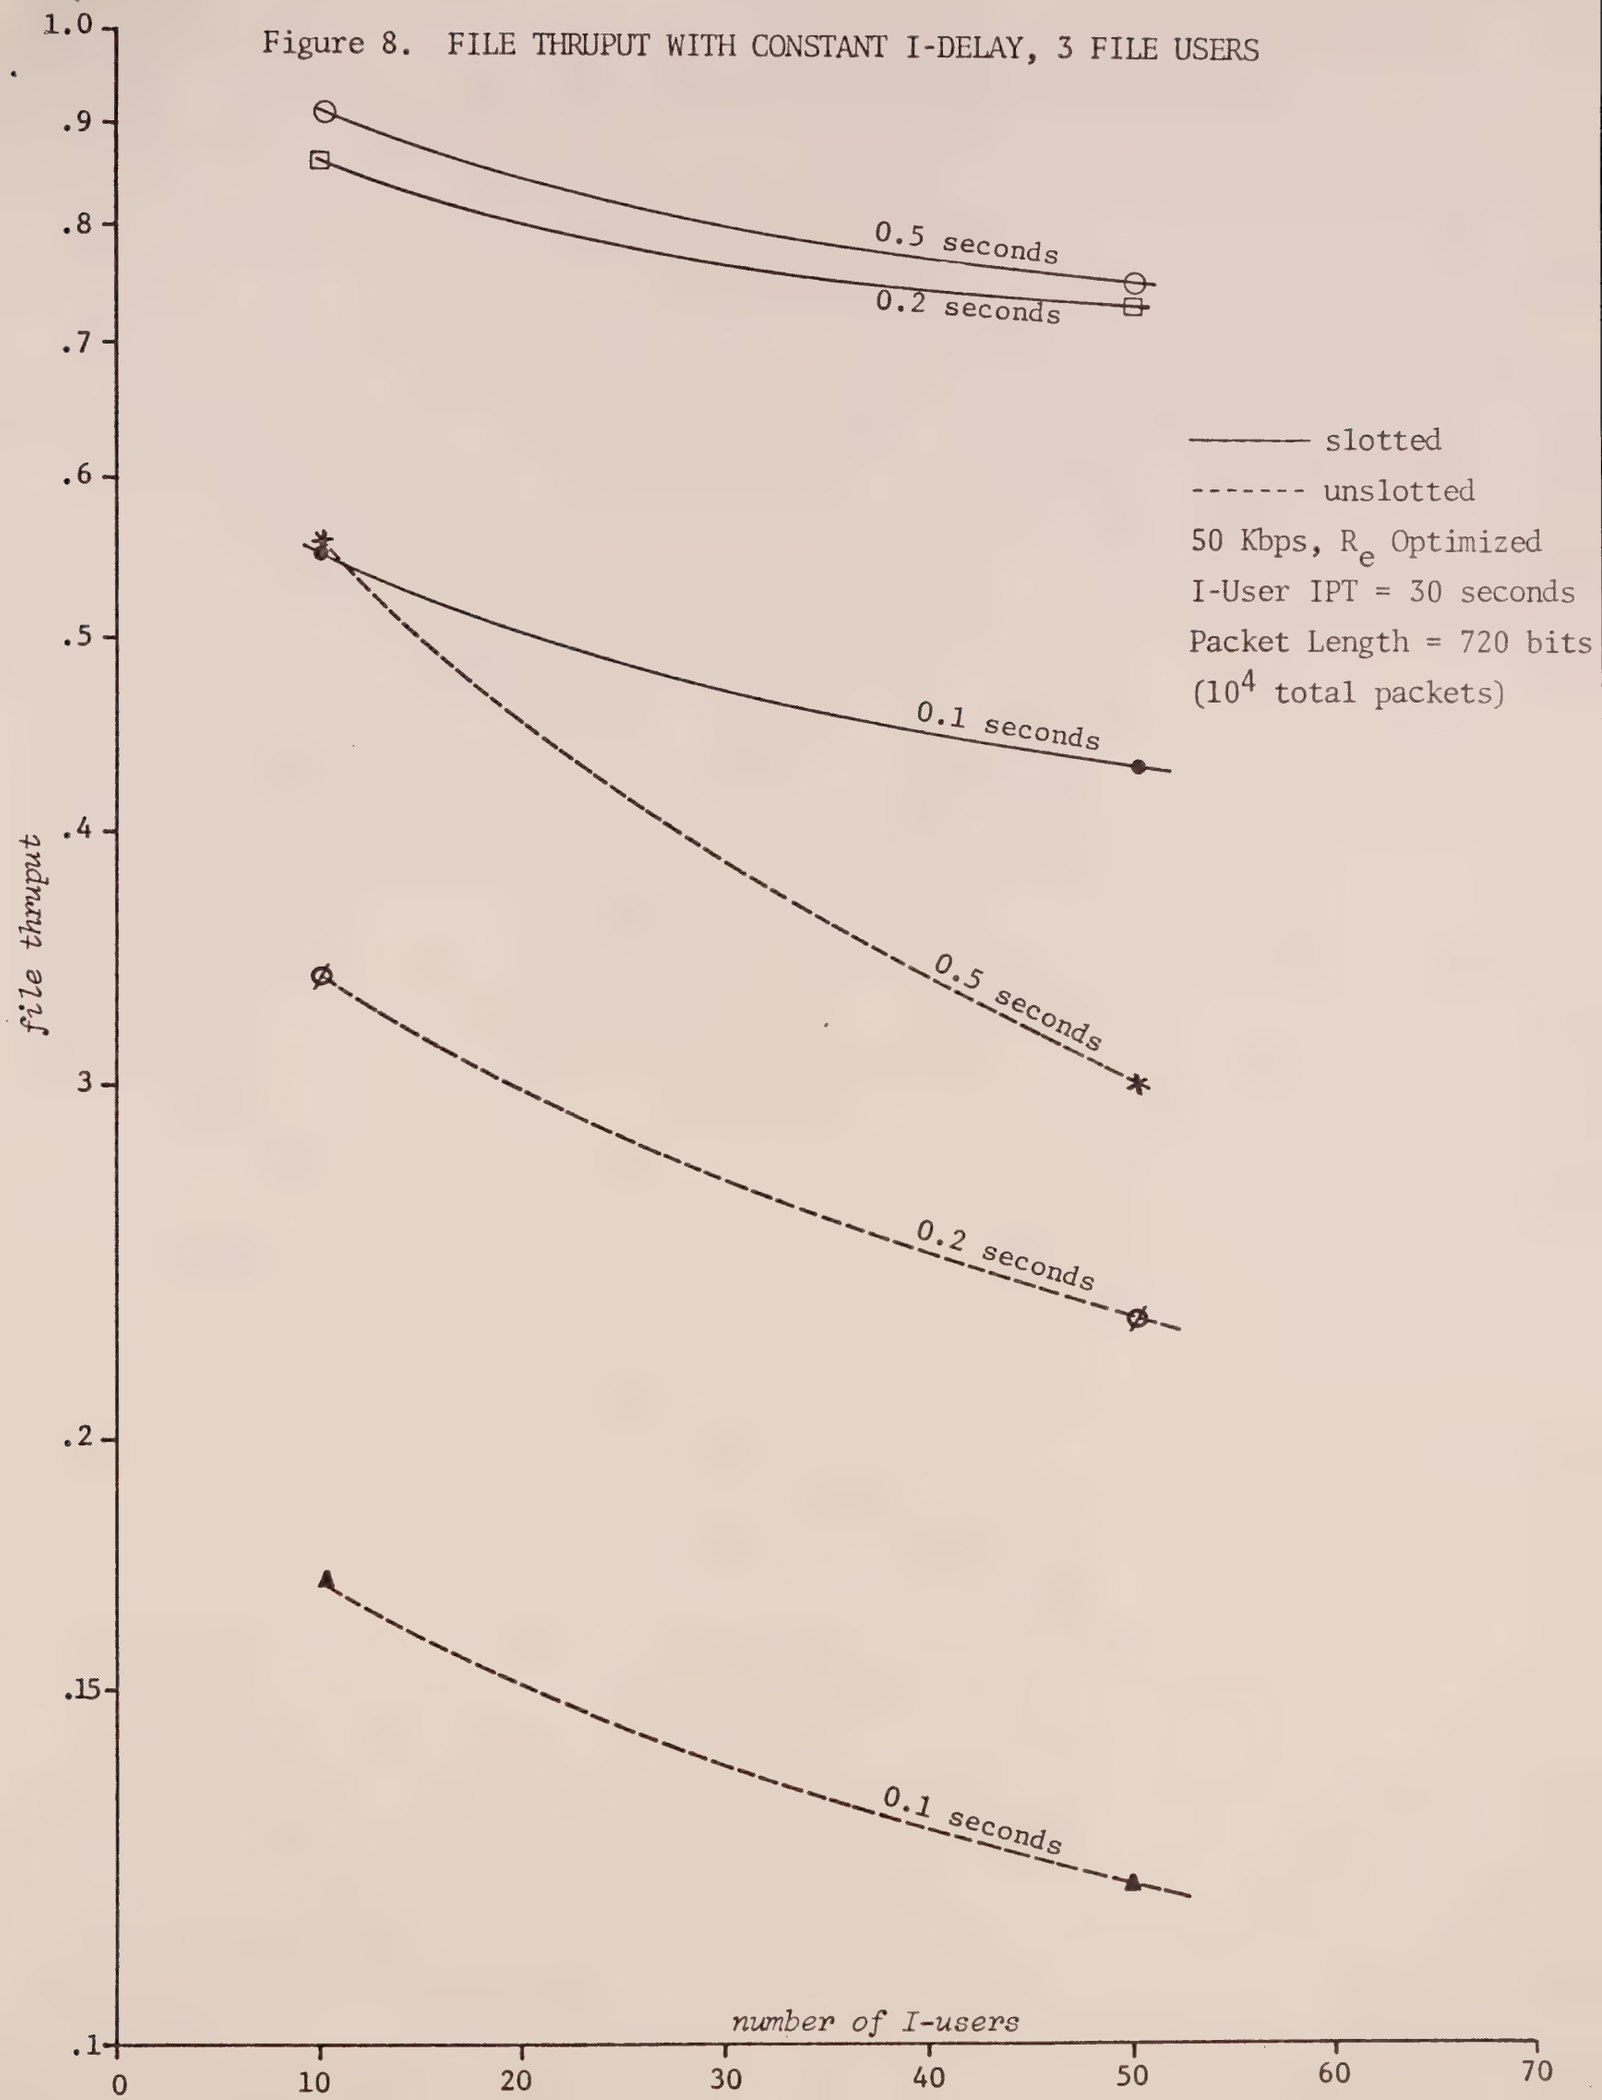

Figure 9. FILE THRUPUT WITH CONSTANT I-DELAY, 10 FILE USERS

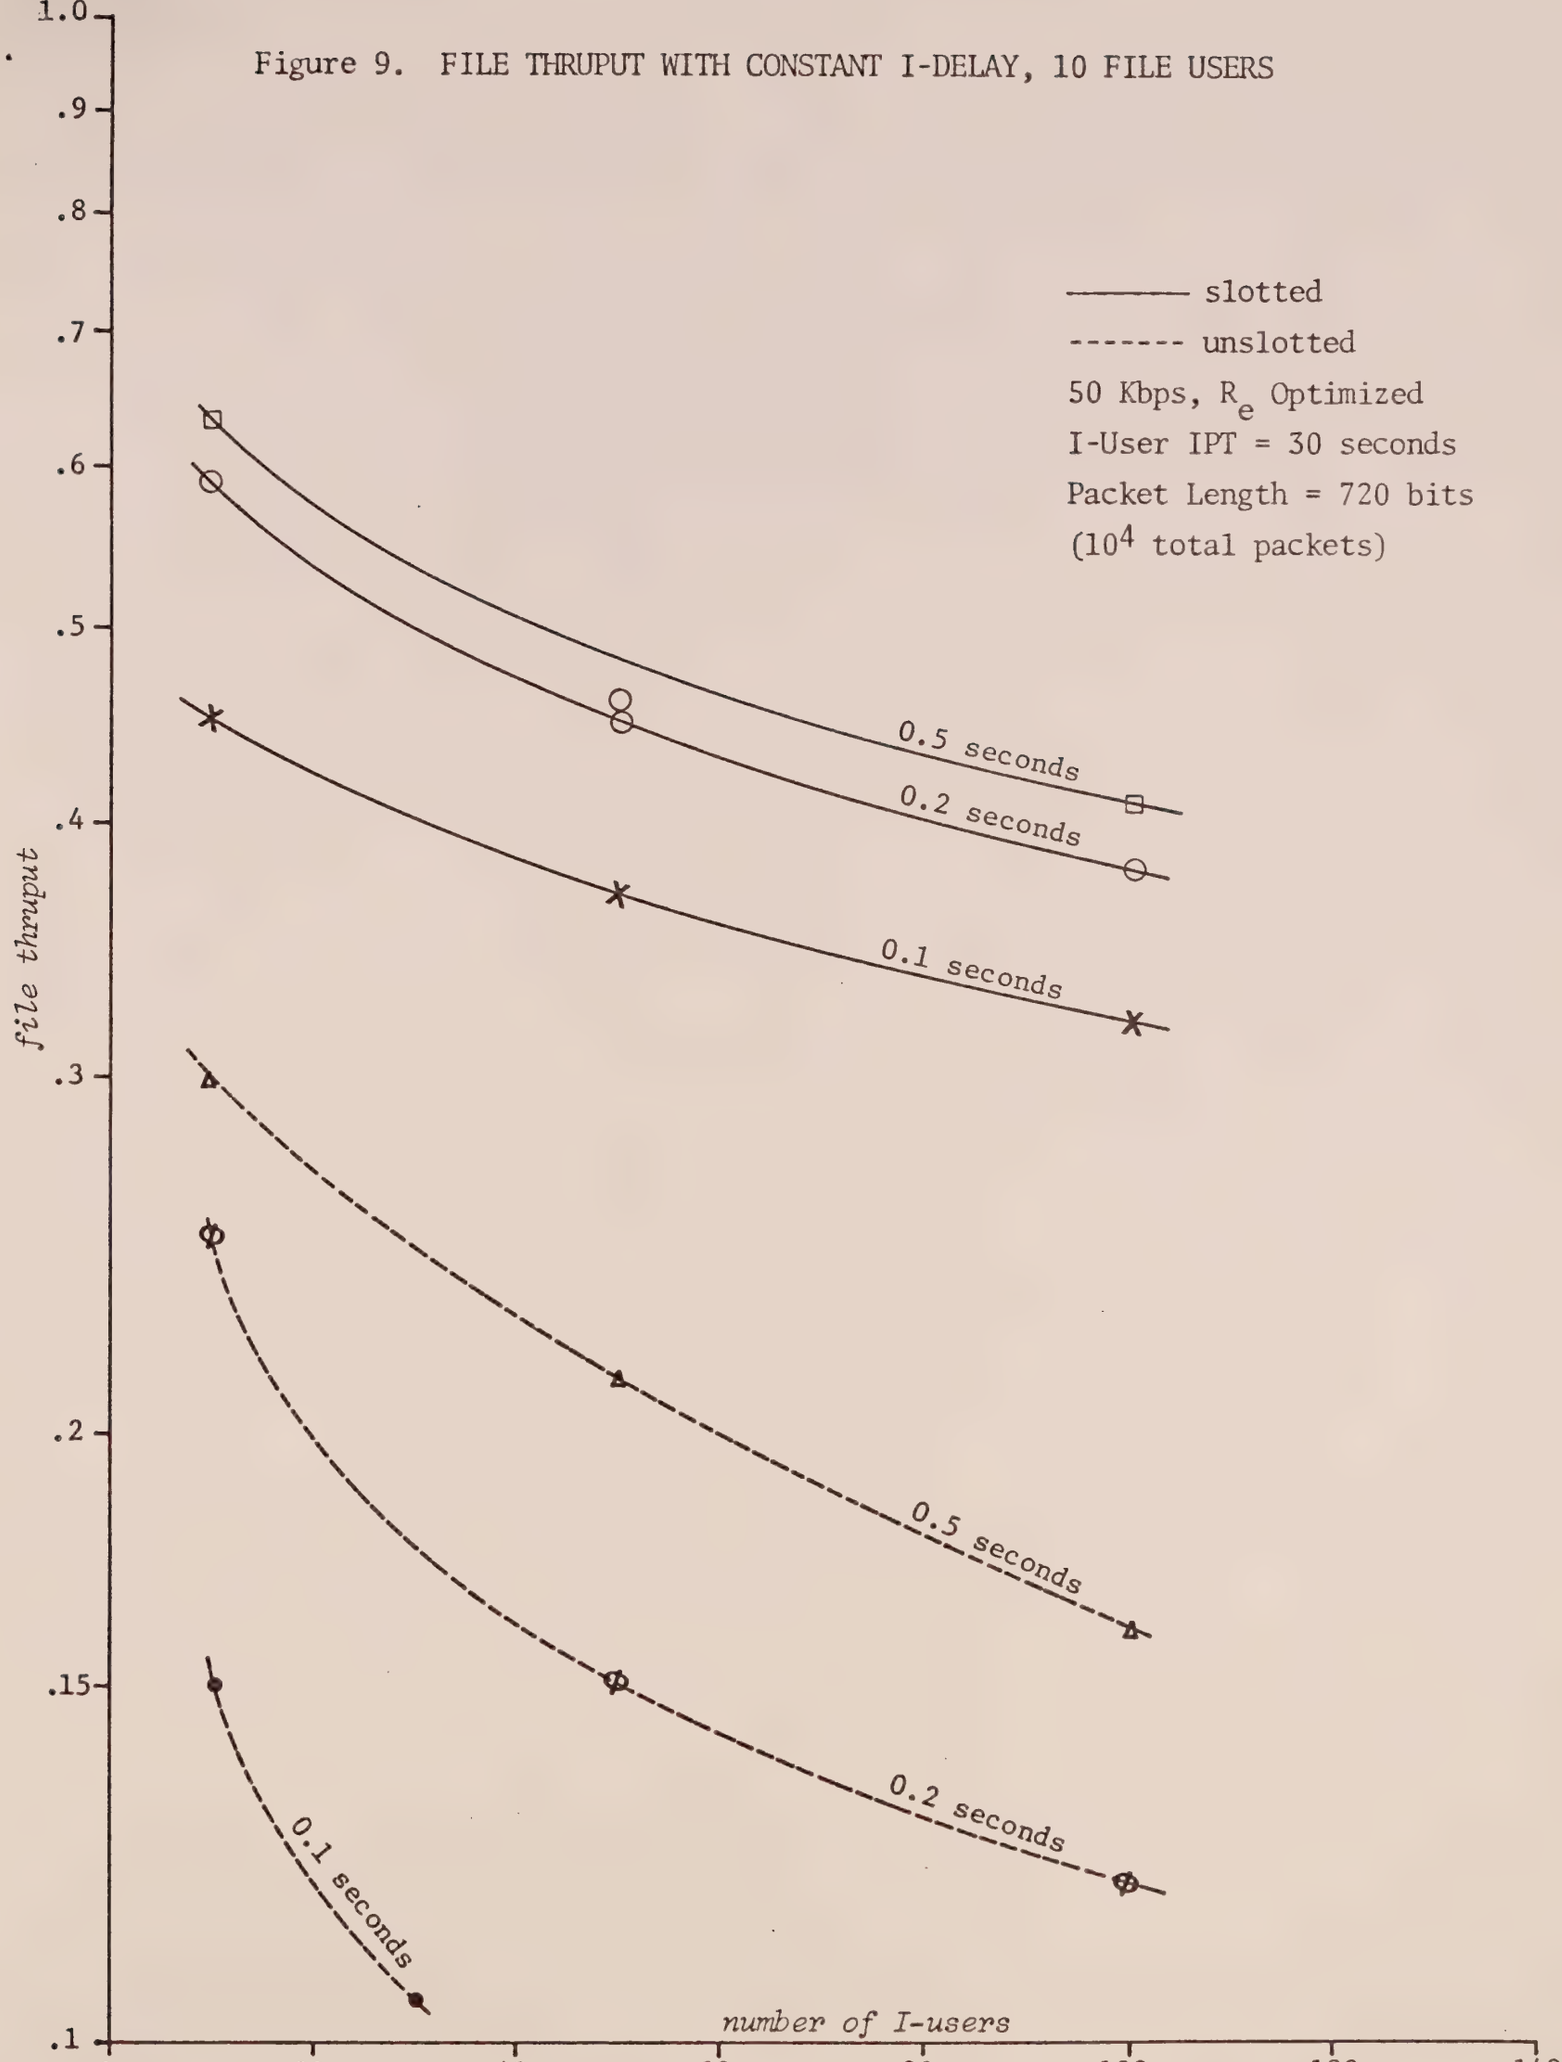

are present -- almost .75 for an I-user delay of 0.2 seconds in the slotted case. Note also the large differences in performance between the slotted and unslotted channels.

Figures 10 and 11 are provided to give the FR value associated with a particular point in Figures 6 to 9. Note the absence of data for the unslotted case, 3 file users, between an FR of .75 and 1.0 in Figure 10; this is due to the limitation of the simulator to integer values of FR. Table 2 gives the optimal values of R_e for the results in this section (also limited to integer values).

Figure 10. FILE THRUPUT VS FRAME RATIO, 3 FILE USERS

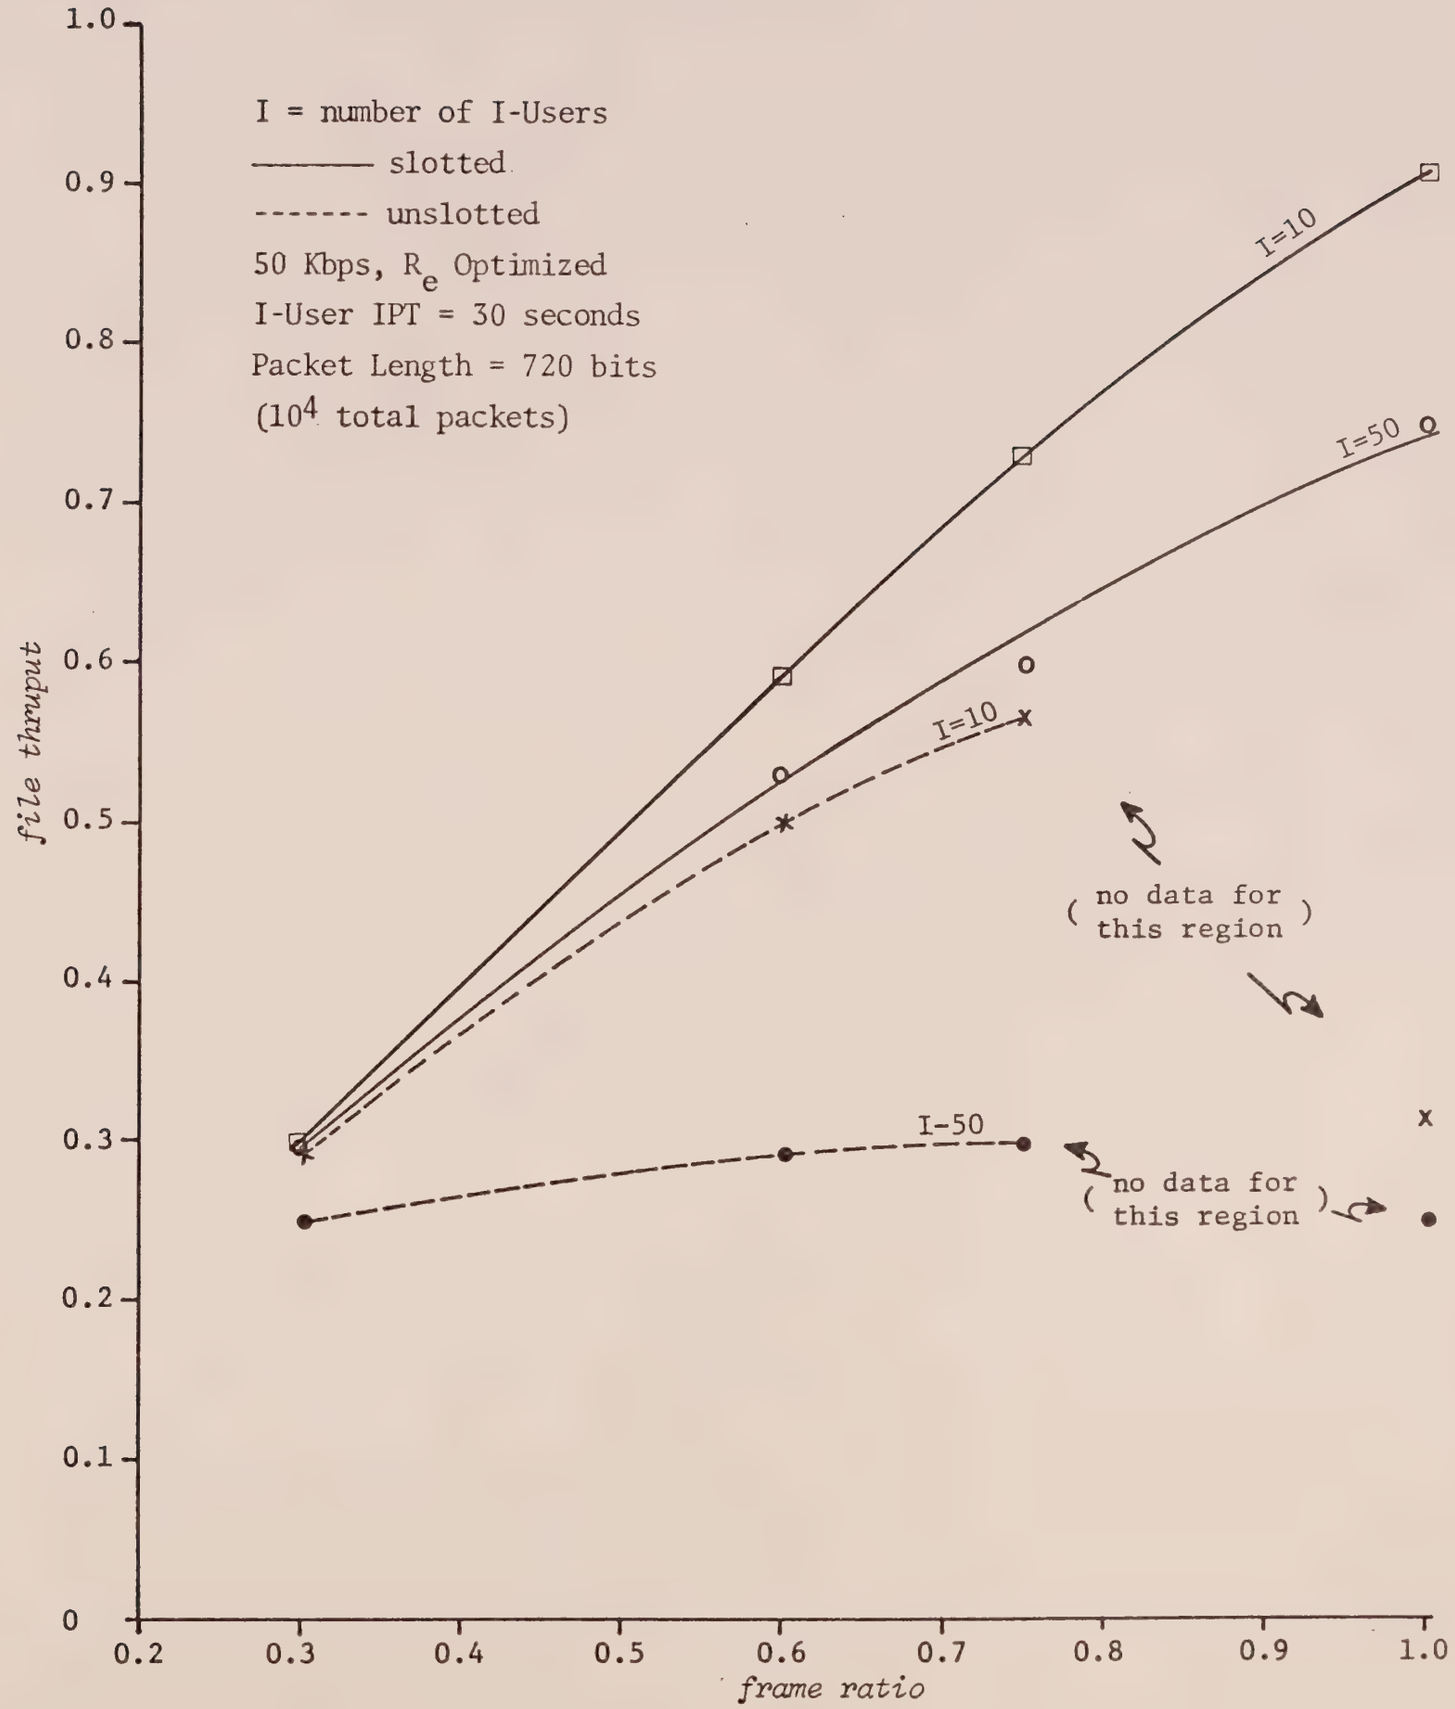

Figure 11. FILE THRUPUT VS FRAME RATIO, 10 FILE USERS

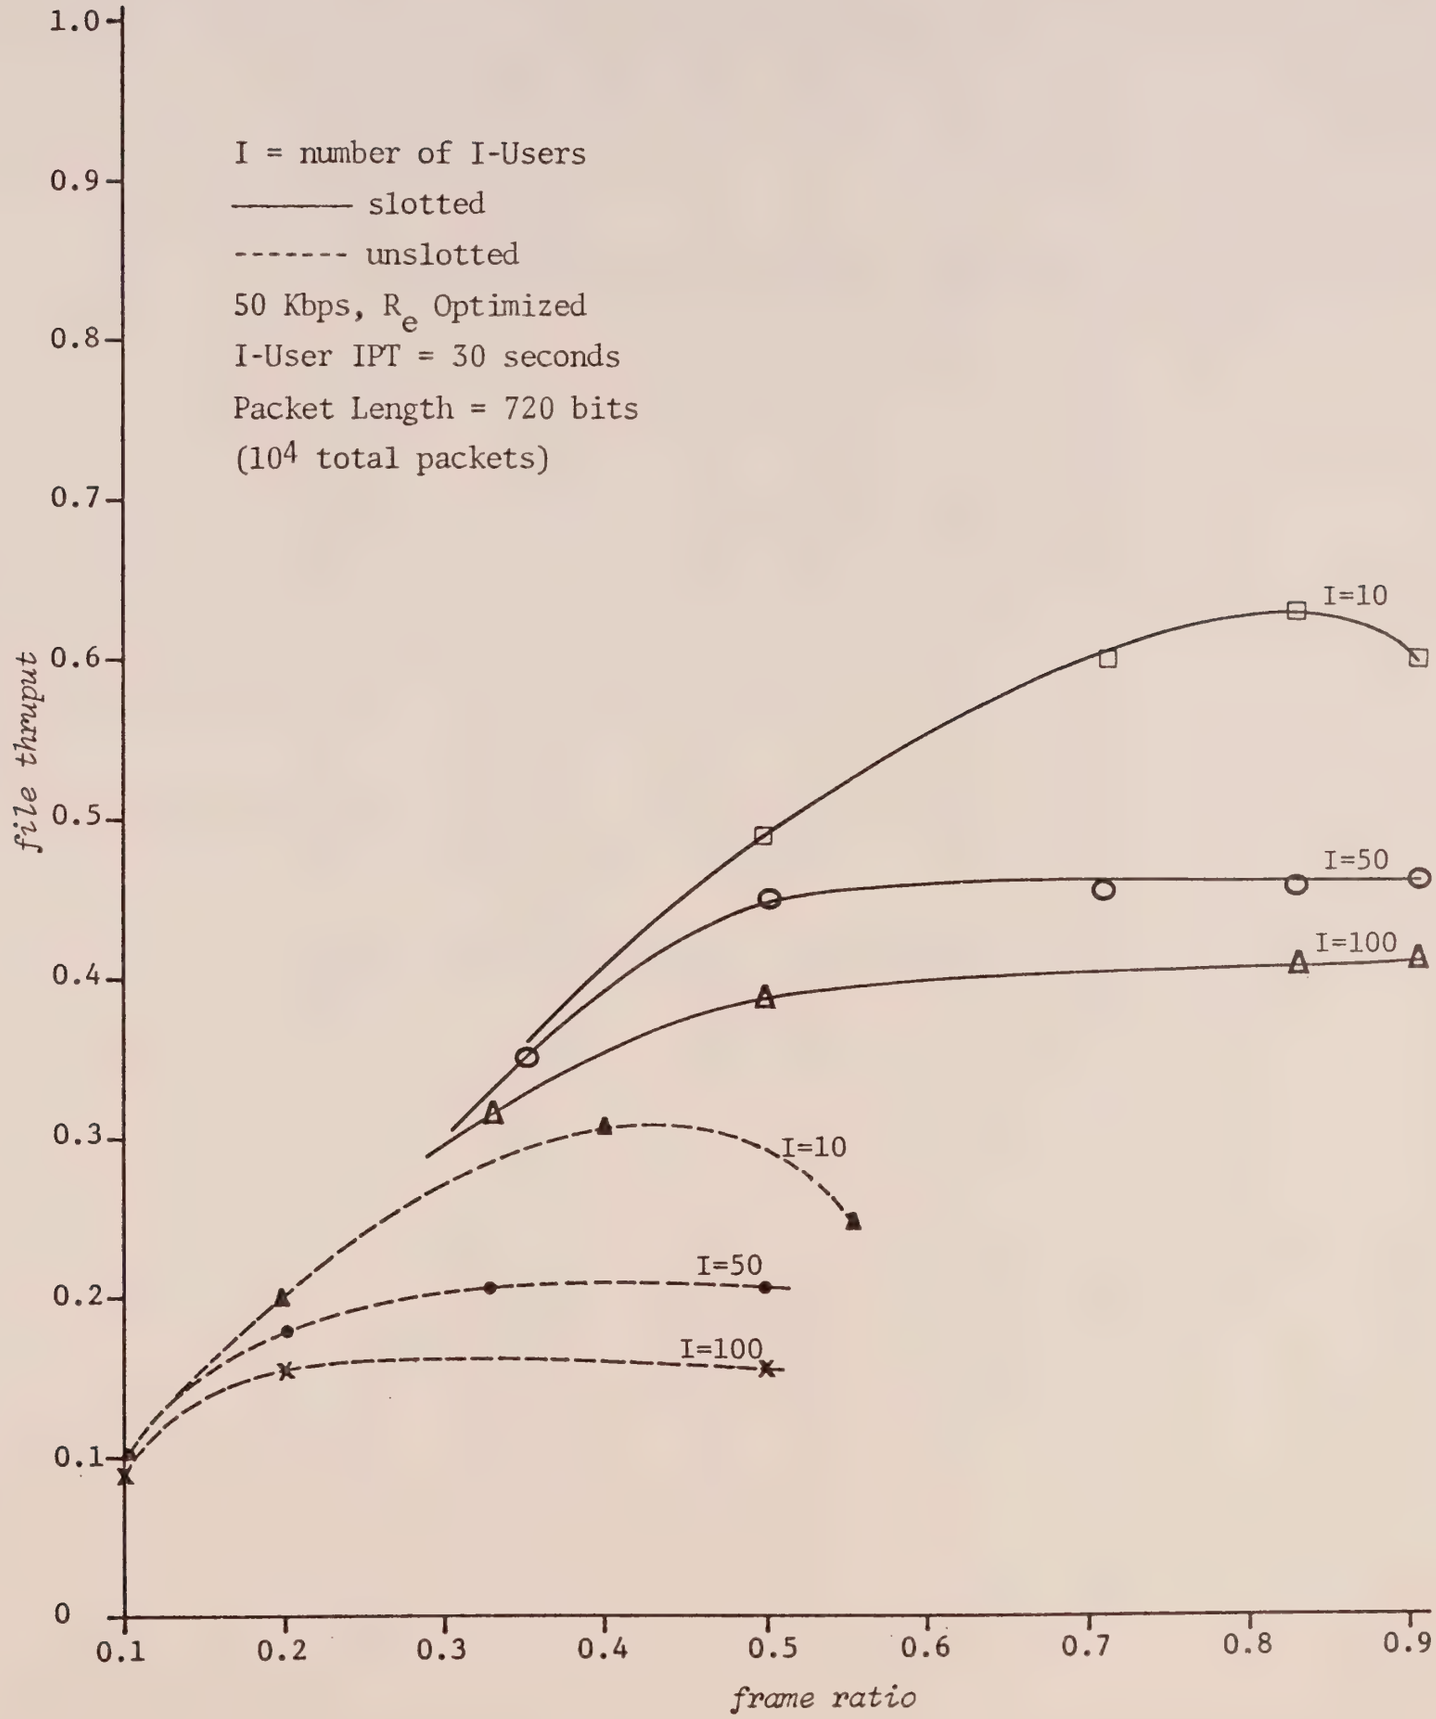

Table 2. OPTIMUM R_e VALUES FOR MIXED TRAFFIC RESULTS

N	I	FR	Optimum R_e (720-bit Packets)	
			slotted	*unslotted*
3	10	3/3	2	20
		3/4	2	2
		3/5	2	3
		3/10	2	5
	50	3/3	3	20
		3/4	3	5
		3/5	3	5
		3/10	3	5
		3/30	-	5
10	10	10/11	15	-
		10/12	5	-
		10/14	5	-
		10/18	-	10
		10/20	4	-
		10/25	-	15
		10/50	4	5
		10/100	-	5
10	50	10/11	10	-
		10/12	10	-
		10/14	5	-
		10/20	4	20
		10/30	4	15
		10/50	4	10
		10/100	-	10
	100	10/11	10	-
		10/12	10	-
		10/20	7	20
		10/30	4	-
		10/50	4	10
		10/100	-	10

5. CONCLUSIONS

This report has presented some simulation results on a technique for regulating file users in a slotted or unslotted ALOHA channel. The technique consists essentially of allowing the file users to automatically synchronize with respect to each other via their random retransmissions, achieved by the provision of sufficiently large 'holes' in the channel between successive new packets sent by each user. By increasing the size of the holes, the delay experienced by interactive users in the channel may be decreased as well.

The method appears to have potential application in a ground-based environment such as the ALOHANET. More study is required of the relative merits of other approaches, as well as the present method in the context of a carrier-sense channel.

REFERENCES

[1] ABRAMSON, NORMAN: "Packet Switching with Satellites," *ALOHA System Technical Report B73-2*, University of Hawaii, March 1973; also published in *AFIPS Conference Proceedings, National Computer Conference*, Vol. 42, New York, June 1973, pp. 695-702.

[2] LAM, SIMON S.: "Packet Switching in a Multi-Access Broadcast Channel with Application to Satellite Communication in a Computer Network," *Report UCLA-ENG-7429*, University of California, School of Engineering and Applied Science, April 1974, pp. 57-74.

[3] CROWTHER, W., et.al.: "A System for Broadcast Communications: Reservation ALOHA," *Proceedings of the Sixth Hawaii International Conference on System Sciences*, Western Periodicals, Inc., January 1973, pp. 371-374.

[4] ROBERTS, L.G.: "Dynamic Allocation of Satellite Capacity through Packet Reservation," *AFIPS Conference Proceedings, National Computer Conference*, Vol. 42, New York, June 1973, pp. 711-716.

[5] BINDER, RICHARD: "A Dynamic Packet-Switching System for Satellite Broadcast Channels," *ALOHA System Technical Report B74-5*, University of Hawaii, August 1974.

[6] KANEHIRA, EARL M.: "Transient and Steady-State Analysis of Reservation ALOHA," *ARPANET Satellite System Note 35, NIC Document 14180*, Stanford Research Institute, Stanford, California, February 1973.

[7] "SIMFAC Version II User's Guide," *ALOHA System General Document CCG/G-74*, University of Hawaii, January 1975.

APPENDIX I: SIMULATION DETAILS

The simulator consists of a program run on a dedicated Lockheed SUE minicomputer. The communications modelled for each I-user closely approximates that of a real system -- whenever a new packet is transmitted, the user is prevented from generating another new packet until an acknowledgement is received. Delay is computed for each new packet sent, and is defined as the time elapsing from the start of transmission of the packet to the end of receipt of an acknowledgement.

I-user 'think-time' is represented by time intervals obtained from an exponential distribution. At the beginning of the simulation, a think-time interval is obtained for each user to represent the start of his first new packet transmission. Subsequently, a new think-time interval is added to the ending time of each acknowledgement packet to obtain the starting time of the next new packet for a user. The same mean think-time is used for all I-users.

Retransmission intervals are also obtained from an exponential distribution, with the same mean (not necessarily the mean think-time) used for all retransmissions by all users. A constant time equal to the transmission time of two maximum length (720-bit) packets is added to each exponential interval to represent a worst-case waiting time for receipt of an acknowledgement by a user. This total is added to the ending time of the last transmission attempt of a user to obtain the starting time of his next retransmission.

All file users are initialized in a blocked state, so that their first transmission will be made after a retransmission waiting interval. Constant time intervals are used between a successful transmission and the next new packet as defined in the main section of this report.